

Proof-relevant π -calculus: a constructive account of concurrency and causality

Roly Perera¹ [†] James Cheney² [‡]

¹ *University of Glasgow; University of Edinburgh*

² *University of Edinburgh*

Received 7 February 2017

We present a formalisation in Agda of the theory of concurrent transitions, residuation, and causal equivalence of traces for the π -calculus. Our formalisation employs de Bruijn indices and dependently-typed syntax, and aligns the “proved transitions” proposed by Boudol and Castellani in the context of CCS with the proof terms naturally present in Agda’s representation of the labelled transition relation. Our main contributions are proofs of the “diamond lemma” for the residuals of concurrent transitions and a formal definition of equivalence of traces up to permutation of transitions.

In the π -calculus transitions represent propagating binders whenever their actions involve bound names. To accommodate these cases, we require a more general diamond lemma where the target states of equivalent traces are no longer identical, but are related by a *braiding* that rewires the bound and free names to reflect the particular interleaving of events involving binders. Our approach may be useful for modelling concurrency in other languages where transitions carry metadata sensitive to particular interleavings, such as dynamically allocated memory addresses.

1. Introduction

The π -calculus [Milner 1999; Milner et al. 1992] is an expressive model of concurrent and mobile processes. It has been investigated extensively and many variants, extensions and refinements proposed, including the asynchronous, polyadic, and applied π -calculus [Sangiorgi and Walker 2001]. The π -calculus has also attracted considerable attention from the logical frameworks and meta-languages community, and formalisations of its syntax and semantics have been developed in most of the extant mechanised metatheory systems, including HOL [Melham 1994; Aït Mohamed 1995], Coq [Hirschhoff 1997; Despeyroux 2000; Honsell et al. 2001], Isabelle/HOL [Röckl et al. 2001; Gay 2001], Isabelle/FM [Gabbay 2003], Nominal Isabelle [Bengtson and Parrow 2009], Abella [Baelde et al. 2014], CLF [Cervesato et al. 2002], and Agda [Orchard and Yoshida 2015]. Indeed, some

[†] Supported by UK EPSRC grant EP/K034413/1 and US AFOSR grant FA8655-13-1-3006.

[‡] Supported by a Royal Society University Research Fellowship.

early formalisations motivated or led to important developments in mechanised metatheory, such as the Theory of Contexts [Bucalo et al. 2006], or CLF’s support for monadic encapsulation of concurrent executions.

Prior formalisations have typically considered the syntax, semantics and bisimulation theory of the π -calculus. One interesting aspect of the π -calculus that has not been formally investigated, and remains to some extent ill-understood informally, is its theory of *causal equivalence*. Two transitions t, t' that can be taken from a process term P are said to be *concurrent*, written $t \smile t'$, if they can be performed “in either order” — that is, if after performing t , there is a natural way to transform the other transition t' so that its effect is performed on the result of t , and vice versa. The transformed version of the transition is said to be the *residual* of t' after t , written t'/t . The key property of this operation, called the “diamond lemma” [Lévy 1980], is that the two residuals t/t' and t'/t result in the same process. Finally, permutation of concurrent transitions induces a *causal equivalence* relation on pairs of traces. This relation is the standard notion of permutation-equivalence from the theory of traces over concurrent alphabets [Mazurkiewicz 1987].

In classical treatments of concurrency and residuation, starting with Lévy [1980], a transition is usually considered to be a triple (e, t, e') where e and e' are the source and target terms of the transition and t is some information about the step performed. Boudol and Castellani [1989] introduced the *proved transitions* approach for CCS in which the labels of transitions are enriched with an approximation of the derivation tree which proves that a particular triple is in the transition relation. Boreale and Sangiorgi [1998] and Degano and Priami [1999] developed theories of causal equivalence for the π -calculus, building indirectly on the proved transition approach; Danos and Krivine [2004] and Cristescu et al. [2013] developed notions of causality in the context of reversible CCS and π -calculus respectively.

None of the above treatments has been mechanised, although the theory of residuals for the λ -calculus was formalised in Coq by Huet [1994] and in Abella by Accattoli [2012]. In this paper, we report on a formalisation of concurrency, residuation and causal equivalence for the π -calculus carried out in the dependently-typed programming language Agda [Norell 2009]. Our approach is inspired by the proved transitions method of Boudol and Castellani. However, by taking a “Church-style” view of the labelled transition semantics and treating transitions as proof terms, rather than triples (e, t, e') , we avoid the need for an auxiliary notion of “proved transition”. Agda’s dependent typing allows us to define the concurrency relation on (compatibly-typed) transition proofs, and residuation as a total function taking two transitions along with a proof that the transitions are concurrent. Our formalisation employs de Bruijn indices [de Bruijn 1972], an approach with well-known strengths and weaknesses compared, for example, to higher-order or nominal abstract syntax techniques employed in existing formalisations; some of these other techniques are discussed in § 5.

Our definition of concurrency is not the only plausible one for the π -calculus. Indeed, there appears to be little consensus regarding the characteristics of a canonical definition. For example, Cristescu et al. [2013] write “[in] the absence of an indisputable definition of permutation equivalence for [labelled transition system] semantics of the π -calculus it is hard to assert the correctness of one definition over another.” We do, however, show that

our definition of concurrency is sound by proving the diamond property; to the best of our knowledge, ours is the first mechanised version of this result for any process calculus.

However, one key observation that emerges in our development is that requiring residuals of concurrent transitions to reach exactly the same state is too restrictive. When the action of a transition involves a bound name, the transition represents a propagating binder. In such cases equivalent traces no longer have identical target states, but rather states which are equal up to a *braiding* that rewires the bound and free names to reflect the different order of events in the two traces. Although typically unobservable to a program, such interleaving-sensitive information may be important for other purposes, such as memory locations in a debugger, or transaction ids in a financial application. In these situations being able to robustly translate between the target states of different interleavings may be important. Our development may therefore be a useful case study for formalising concurrency in other settings where transition labels carry interleaving-sensitive metadata.

This is a substantially revised version of a paper presented at the *Logical Frameworks and Meta-Languages: Theory and Practice* workshop [Perera and Cheney 2015]. This version extends the earlier work with graphical proof-sketches for various lemmas, a more detailed comparison of related formalisation efforts, extensive examples and discussion regarding the generalised diamond property, a more precise definition of cofinality, and a formalisation of composite braids. A companion paper [Perera et al. 2016] uses the formalisation of concurrent transitions presented here as the basis for “causally consistent” dynamic slicing of π -calculus programs.

The paper is organised as follows. §2 describes our variant of the (synchronous) π -calculus, including syntax, renamings, and transitions. §3 defines concurrency and residualisation for transitions, and discusses the diamond lemma and the notion of “cofinal” transitions. §4 presents our definition of causal equivalence. §5 discusses related work in more detail and §6 concludes and discusses prospects for future work. Appendix A summarises the Agda module structure; the source code can be found at <https://github.com/rolyp/proof-relevant-pi>, release 0.3.

2. Synchronous π -calculus

We present our formalisation in the setting of a first-order, synchronous, monadic π -calculus with recursion and internal choice, using a labelled transition semantics.

Names are ranged over by x, y and z . An input action is written \underline{x} . Output actions are written $\bar{x}(y)$ if y is in scope and \bar{x} if the action represents the output of a name whose scope is extruding, in which case we say the action is a *bound* output. Bound outputs do not appear in source programs but arise during execution.

Name	$x, y, z ::= 0 \mid 1 \mid \dots$	
Action	$a ::= \underline{x}$	input
	$\bar{x}(y)$	output
	\bar{x}	bound output
	τ	silent

Although it has become common practice to limit attention to sums of guarded processes, here we study the calculus as originally formulated by Milner et al, which allows sums of arbitrary processes. (Our basic approach should transfer to guarded choice, and other common variants.)

Process	$P, Q, R, S ::= \mathbf{0}$	inactive
	$\underline{x}.P$	input
	$\bar{x}(y).P$	output
	$P + Q$	non-guarded choice
	$P \mid Q$	parallel
	νP	restriction
	$!P$	replication

Although the formal development uses de Bruijn indices, and we give definitions and state properties in terms of this notation, we will sometimes illustrate their meaning in terms of conventional π -calculus notation. For example, the conventional π -calculus term $(\nu x) x(z).\bar{y}(z).\mathbf{0} \mid \bar{x}(c).\mathbf{0}$ would be represented using de Bruijn indices as $\nu(\underline{0}.n+1\langle 0 \rangle).\mathbf{0} \mid \bar{0}(m+1).\mathbf{0}$, provided that y and c are associated with indices n and m . Here, the first $\mathbf{0}$ represents the bound variable x , the second $\mathbf{0}$ the bound variable z , and the third refers to x again. Note that the symbol $\mathbf{0}$ denotes the inactive process, not a de Bruijn index.

The syntax of actions and processes is defined more formally in Figure 1 overleaf. Let Γ and Δ range over *contexts*, which in an untyped setting are simply natural numbers. A membership witness $x \in \Gamma$ is a proof that $x < \Gamma$. A context Γ *closes* P if $x \in \Gamma$ for every free variable x of P . We denote by $\text{Proc } \Gamma$ the set of processes closed by Γ , as defined below. We write $\Gamma \vdash P$ to mean $P \in \text{Proc } \Gamma$. Similarly, actions are well-formed only in closing contexts; we write $a : \text{Action } \Gamma$ to mean that Γ is closing for a .

To specify the labelled transition semantics, it is convenient to distinguish *bound* actions b from non-bound actions c . A bound action $b : \text{Action } \Gamma$ is of the form \underline{x} or \bar{x} , and shifts a process from Γ to a target context $\Gamma + 1$, freeing the index $\mathbf{0}$. A non-bound action $c : \text{Action } \Gamma$ is of the form $\bar{x}(y)$ or τ , and has a target context which is also Γ . Meta-variable a ranges over all actions, bound and non-bound. $|a|$ denotes the amount by which the action increments the context; thus $|b| = 1$ and $|c| = 0$.

2.1. Renamings

A de Bruijn indices formulation of π -calculus makes extensive use of renamings. A *renaming* $\rho : \Gamma \longrightarrow \Delta$ is any function (injective or otherwise) from names in Γ to names in Δ . The labelled transition semantics makes use of the lifting of the successor function

$$\begin{array}{c}
\boxed{\Gamma \vdash P} \\
\\
\frac{}{\Gamma \vdash 0} \quad \frac{\Gamma + 1 \vdash P}{\Gamma \vdash \underline{x}.P} \quad x \in \Gamma \quad \frac{\Gamma \vdash P}{\Gamma \vdash \bar{x}(y).P} \quad x, y \in \Gamma \quad \frac{\Gamma \vdash P \quad \Gamma \vdash Q}{\Gamma \vdash P + Q} \\
\\
\frac{\Gamma \vdash P \quad \Gamma \vdash Q}{\Gamma \vdash P | Q} \quad \frac{\Gamma + 1 \vdash P}{\Gamma \vdash \nu P} \quad \frac{\Gamma \vdash P}{\Gamma \vdash !P} \\
\\
\boxed{a : \text{Action } \Gamma} \\
\\
\frac{}{\underline{x} : \text{Action } \Gamma} \quad x \in \Gamma \quad \frac{}{\bar{x} : \text{Action } \Gamma} \quad x \in \Gamma \quad \frac{}{\bar{x}(y) : \text{Action } \Gamma} \quad x, y \in \Gamma \quad \frac{}{\tau : \text{Action } \Gamma}
\end{array}$$

Fig. 1: Syntax of processes and actions

$\cdot + 1$ on natural numbers to renamings, which we call **push** to avoid confusion with the $\cdot + 1$ operation on contexts; **pop** y , which undoes the effect of **push**, replacing 0 by y ; and **swap**, which transposes the roles of 0 and 1 but otherwise acts as the identity. This de Bruijn treatment of π -calculus is similar to that of Hirschhoff's asynchronous μ s calculus [Hirschhoff 1999]; in particular Hirschhoff's $\langle x \rangle$, ϕ and ψ operators correspond roughly to **pop** x , **push** and **swap**. We give a late rather than early semantics; other differences are discussed in § 5 below.

$$\begin{array}{ccc}
\boxed{\text{push}_\Gamma : \Gamma \longrightarrow \Gamma + 1} & \boxed{\text{pop}_\Gamma y : \Gamma + 1 \longrightarrow \Gamma} & \boxed{\text{swap}_\Gamma : \Gamma + 2 \longrightarrow \Gamma + 2} \\
\\
\text{push } x = x + 1 & \text{pop } y \ 0 = y \\
& \text{pop } y \ (x + 1) = x & \text{swap } 0 = 1 \\
& & \text{swap } 1 = 0 \\
& & \text{swap } (x + 2) = x + 2
\end{array}$$

Fig. 2: **push**, **pop** and **swap** renamings

The Γ subscripts that appear on push_Γ , $\text{pop}_\Gamma y$ and swap_Γ are shown in grey to indicate that they may be omitted when their value is obvious or irrelevant; this is a convention we use throughout the paper.

2.1.1. Lifting renamings to processes and actions The functorial extension $\rho^* : \text{Proc } \Gamma \longrightarrow \text{Proc } \Delta$ of a renaming $\rho : \Gamma \longrightarrow \Delta$ to processes is defined in the usual way. Renaming under a binder utilises the action of $\cdot + 1$ on renamings, which is also functorial. Syntactically, ρ^* binds tighter than any process constructor, and $\cdot + 1$ has higher precedence than composition, so that (for example) $\text{pop } 0 \circ \text{push} + 1$ means $\text{pop } 0 \circ (\text{push} + 1)$, not $(\text{pop } 0 \circ \text{push}) + 1$.


<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> $\cdot^* : (\Gamma \longrightarrow \Delta) \longrightarrow \text{Proc } \Gamma \longrightarrow \text{Proc } \Delta$ </div> $ \begin{aligned} \rho^* \mathbf{0} &= \mathbf{0} \\ \rho^* (\underline{x}.P) &= \underline{\rho x} . (\rho + 1)^* P \\ \rho^* (\overline{x}\langle y \rangle . P) &= \overline{\rho x} \langle \rho y \rangle . \rho^* P \\ \rho^* (P + Q) &= \rho^* P + \rho^* Q \\ \rho^* (P \mid Q) &= \rho^* P \mid \rho^* Q \\ \rho^* (\nu P) &= \nu (\rho + 1)^* P \\ \rho^* (!P) &= !\rho^* P \end{aligned} $	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> $\cdot^* : (\Gamma \longrightarrow \Delta) \longrightarrow \text{Action } \Gamma \longrightarrow \text{Action } \Delta$ </div> $ \begin{aligned} \rho^* \underline{x} &= \underline{\rho x} \\ \rho^* \overline{x} &= \overline{\rho x} \\ \rho^* \tau &= \tau \\ \rho^* \overline{x}\langle y \rangle &= \overline{\rho x} \langle \rho y \rangle \end{aligned} $
<div style="border: 1px solid black; padding: 5px; margin: 0 auto; width: 60%;"> $\cdot + 1 : (\Gamma \longrightarrow \Delta) \longrightarrow \Gamma + 1 \longrightarrow \Delta + 1$ </div> $ \begin{aligned} (\rho + 1) \mathbf{0} &= \mathbf{0} \\ (\rho + 1) (x + 1) &= \rho x + 1 \end{aligned} $	

Fig. 3: Renaming for processes and actions


2.1.2. Properties of renamings Several equational properties of renamings are used throughout the development; here we present the ones mentioned elsewhere in the paper. For each lemma, we give the corresponding commutative diagram underneath on the left, along with a string diagram that offers a graphical intuition for why the lemma holds.

Lemma 1. $\text{pop } x \circ \text{push} = \text{id}$


Freeing the index 0 and then immediately substituting x for it is a no-op.



Lemma 2. $\text{pop } 0 \circ \text{push} + 1 = \text{id}$




Lemma 3. $\text{swap} + 1 \circ \text{swap} \circ \text{swap} + 1 = \text{swap} \circ \text{swap} + 1 \circ \text{swap}$




The above are two ways to swap indices 0 and 2.

Lemma 4. $\text{pop } 0 \circ \text{swap} = \text{pop } 0$



Lemma 5. $\text{swap} \circ \text{push} + 1 = \text{push}$, $\text{swap} \circ \text{push} = \text{push} + 1$




Lemma 6. $\text{push} \circ \rho = \rho + 1 \circ \text{push}$

Lemma 7. $\rho \circ \text{pop } x = \text{pop } \rho x \circ \rho + 1$

Lemma 8. $\text{swap} \circ \rho + 2 = \rho + 2 \circ \text{swap}$

These last three lemmas assert various naturality properties of push , $\text{pop } x$ and swap .



2.2. Labelled transition semantics

An important feature of our presentation is that each transition rule has an explicit constructor name. This allow derivations to be written in a compact, expression-like form, similar to the *proven transitions* used by Boudol and Castellani [1989] to define notions of concurrency and residuation for CCS. However, rather than giving an additional inductive definition describing the structure of a “proof” that $P \xrightarrow{a} R$, we simply treat the inductive definition of \xrightarrow{a} as a data type. This is a natural approach in a dependently-typed setting.

The rule names are summarised below, and have been chosen to reflect, where possible, the structure of the process triggering the rule. The corresponding relation $P \xrightarrow{a} R$ is defined in Figure 4, for any process $\Gamma \vdash P$, any $a : \text{Action } \Gamma$ and any $\Gamma + |a| \vdash R$.

Transition	$t, u ::=$	$\underline{x}.P$	input on x
		$\overline{x}(y).P$	output y on x
		$t + Q$	choose left or right branch
		$t \mid Q$	propagate a through parallel composition on the left or right
		$t \mid_y u$	synchronise (receiving y on the left or right)
		$\overline{v}t$	initiate extrusion of v
		$t \mid_v u$	v -synchronise (receiving 0 on the left or right)
		$v^a t$	propagate a through binder
		$!t$	replicate

$P \xrightarrow{a} R$

$$\begin{array}{c}
\underline{x}.P \xrightarrow{\underline{x}} \underline{x}.P \quad \overline{x}(y).P \xrightarrow{\overline{x}(y)} \overline{x}(y).P \quad \cdot + Q \xrightarrow{a} P + Q \xrightarrow{a} R \\
\cdot \mid Q \xrightarrow{c} P \mid Q \xrightarrow{c} R \mid Q \quad \cdot \mid Q \xrightarrow{b} P \mid Q \xrightarrow{b} R \mid \text{push}^* Q \\
\cdot \mid_y \cdot \xrightarrow{\underline{x}} P \mid Q \xrightarrow{\tau} (\text{pop } y)^* R \mid S \quad \overline{v}. \xrightarrow{\overline{x+1}(0)} P \xrightarrow{\overline{x}} vP \xrightarrow{\overline{x}} R \\
\cdot \mid_v \cdot \xrightarrow{\underline{x}} P \mid Q \xrightarrow{\tau} v(R \mid S) \quad v^c. \xrightarrow{\text{push}^* c} P \xrightarrow{c} vP \xrightarrow{c} vR \\
v^b. \xrightarrow{\text{push}^* b} P \xrightarrow{b} vP \xrightarrow{b} v(\text{swap}^* R) \quad !. \xrightarrow{a} P \mid !P \xrightarrow{a} R \quad !P \xrightarrow{a} R
\end{array}$$

Fig. 4: Labelled transition rules ($P + \cdot$, $P \mid^b \cdot$, $P \mid^c \cdot$, $\cdot \mid_v \cdot$ and $\cdot \mid_y \cdot$ variants omitted)

The constructor name for each rule is shown to the left of the rule. There is an argument

position, indicated by \cdot , for each premise of the rule. Note that there are two forms of the transition constructors $\cdot^a| \cdot$ and $\nu^a \cdot$ distinguished by whether they are indexed by a bound action b or by a non-bound action c . Omitted from Figure 4 are additional (but symmetric) rules of the form $P + \cdot$, $P| \cdot^b$ and $P| \cdot^b$ where the sub-transition occurs on the opposite side of the operator, and also $\cdot_y| \cdot$ (synchronise) and $\cdot_\nu| \cdot$ (ν -synchronise) rules in which the positions of sender and receiver are transposed. These are all straightforward variants of the rules shown, and are omitted from the figure to avoid clutter. Meta-variables t and u range over transition derivations; if $t : P \xrightarrow{a} R$ then $\text{src}(t)$ denotes P and $\text{tgt}(t)$ denotes R .

Although a de Bruijn formulation of π -calculus requires a certain amount of house-keeping, one pleasing consequence is that the usual side-conditions associated with the π -calculus transition rules are either subsumed by syntactic constraints on actions, or “operationalised” using the renamings above. In particular:

- 1 The use of **push** in the $\cdot^b| Q$ rule corresponds to the usual side-condition asserting that the binder being propagated by P is not free in Q . In the de Bruijn setting every binder “locally” has the name 0, and so this requirement can be operationalised by rewiring Q so that the name 0 is reserved. The **push** will be matched by a later **pop** which substitutes for 0, in the event that the action has a successful synchronisation.
- 2 The $\bar{\nu} \cdot$ rule requires an extrusion to be initiated by an output of the form $\bar{x} + 1\langle 0 \rangle$, capturing the usual side-condition that the name being extruded *on* is distinct from the name being extruded.
- 3 The rules of the form $\nu^a \cdot$ require that the action being propagated has the form **push** *a , ensuring that it contains no uses of index 0. This corresponds to the usual requirement that an action can only propagate through a binder that it does not mention.

The use of **swap** in the ν^b case follows Hirschhoff [1999] and has no counterpart outside of the de Bruijn setting. As a propagating binder passes through another binder, their local indices are 0 and 1. Propagation transposes the binders, and so to preserve naming we rewire R with a “braid” that swaps 0 and 1. Since binders are also reordered by *permutations* that relate causally equivalent executions, the **swap** renaming will also play an important role when we consider concurrent transitions (§3).

The following schematic derivation shows how the compact notation works. Suppose $t : P \xrightarrow{z+2(0)} R$ takes place immediately under a ν -binder, causing the scope of the binder to be extruded. Then suppose the resulting bound output propagates through another binder, giving the partial derivation on the left:

$$\begin{array}{c}
 \vdots \\
 \textcolor{blue}{t} \frac{}{P \xrightarrow{z+2(0)} R} \\
 \textcolor{blue}{\bar{\nu}} \cdot \frac{}{\nu P \xrightarrow{z+1} R} \\
 \textcolor{blue}{\nu^z} \cdot \frac{}{\nu \nu P \xrightarrow{z} \nu R}
 \end{array}
 \qquad
 \begin{array}{c}
 \vdots \\
 \textcolor{blue}{\bar{\nu}t} \frac{}{\nu P \xrightarrow{z+1} R} \\
 \textcolor{blue}{\nu^z} \cdot \frac{}{\nu \nu P \xrightarrow{z} \nu R}
 \end{array}
 \qquad
 \begin{array}{c}
 \vdots \\
 \textcolor{blue}{\nu^z \bar{\nu}t} \frac{}{\nu \nu P \xrightarrow{z} \nu R}
 \end{array}$$

with t standing in for the rest of the derivation. The constructors annotating the left-hand side of the derivation tree (shown in blue in the electronic version of this article) can be thought of as a partially unrolled “transition term” representing the proof. The \cdot placeholders associated with each constructor are conceptually filled by the transition terms annotating the premises of that step. We can “roll up” the derivation by a single

step, by moving the premises into their corresponding placeholders, as shown in the middle figure.

By repeating this process, we can write the whole derivation compactly as $\nu^{\overline{z}} \overline{\nu} t$, as shown on the right. Thus the compact form is simply a flattened transition derivation: similar to a simply-typed λ -calculus term written as a conventional expression, in a (Church-style) setting where a term is, strictly speaking, a typing derivation.

2.2.1. Residuals of transitions and renamings A transition t with action a survives any suitably-typed renaming ρ . Moreover ρ has an image in t , which is simply $\rho + |a|$.

Lemma 9. Suppose $t : P \xrightarrow{a} Q$ and $\rho : \Gamma \longrightarrow \Delta$, where $\Gamma \vdash P$. Then there exists a transition $\rho^* t : \rho^* P \xrightarrow{\rho^* a} (\rho + |a|)^* Q$.

$$\begin{array}{ccc} P & \xrightarrow{t^c} & Q \\ \rho^* \downarrow & & \downarrow \rho^* \\ \rho^* P & \xrightarrow{(\rho^* t)^{\rho^* c}} & \rho^* Q \end{array} \quad \begin{array}{ccc} P & \xrightarrow{t^b} & Q \\ \rho^* \downarrow & & \downarrow (\rho + 1)^* \\ \rho^* P & \xrightarrow{(\rho^* t)^{\rho^* b}} & (\rho + 1)^* Q \end{array}$$

Proof. By the following defining equations. The various renaming lemmas needed to enable the induction hypothesis in each case are omitted.

$\rho^* t^c$

$$\begin{aligned} \rho^*(\overline{x}\langle y \rangle.P) &= \overline{\rho x}\langle \rho y \rangle.\rho^* P \\ \rho^*(t + Q) &= \rho^* t + \rho^* Q \\ \rho^*(P + u) &= \rho^* P + \rho^* u \\ \rho^*(P \mid^c u) &= \rho^* P \mid^{\rho^* c} \rho^* u \\ \rho^*(t \mid^c Q) &= \rho^* t \mid^{\rho^* c} \rho^* Q \\ \rho^*(t \mid_y u) &= \rho^* t \mid_{\rho^* y} \rho^* u \\ \rho^*(t \mid_v u) &= \rho^* t \mid_v \rho^* u \\ \rho^*(\nu^c t) &= \nu^{\rho^* c} (\rho + 1)^* t \\ \rho^*(!t) &= !\rho^* t \end{aligned}$$

$\rho^* t^b$

$$\begin{aligned} \rho^*(\underline{x}.P) &= \underline{\rho x}.\rho^* P \\ \rho^*(t + Q) &= \rho^* t + \rho^* Q \\ \rho^*(P + u) &= \rho^* P + \rho^* u \\ \rho^*(P \mid^b u) &= \rho^* P \mid^{\rho^* b} \rho^* u \\ \rho^*(t \mid^b Q) &= \rho^* t \mid^{\rho^* b} \rho^* Q \\ \rho^*(\overline{\nu} t) &= \overline{\nu} (\rho + 1)^* t \\ \rho^*(\nu^b t) &= \nu^{\rho^* b} (\rho + 1)^* t \\ \rho^*(!t) &= !\rho^* t \end{aligned}$$

We would not expect $\rho^* t$ to be a derivable transition, and thus Lemma 9 to hold, for arbitrary ρ in all extensions of the π -calculus. In particular, the mismatch operator $[x \neq y]P$ that steps to P if x and y are distinct names is only stable under injective renamings.

2.2.2. Structural congruences Our LTS semantics is standard and therefore closed under the usual π -calculus congruences. Structural congruences can be formalised as a bisimulation, using an analogue of the notion of residuation with respect to a transition used elsewhere in this paper. This remains out of scope of the present development.

3. Concurrency, residuals and cofinality

Transitions $P \xrightarrow{a} R$ and $Q \xrightarrow{a'} S$ are *coinitial* when $P = Q$. In this section we formalise a symmetric, irreflexive *concurrency* relation \smile over coinitial transitions. Concurrent transitions $t \smile t'$ are independent, or causally unordered. In an interleaving semantics, t and t' can execute in either order without significant interference; in a true concurrency setting, t and t' form a single, two-dimensional “parallel move” [Curry and Feys 1958]. Concurrency was explored notably by Lévy [1980] for the λ -calculus, and later by Stark [1989] for arbitrary transition systems. The inspiration for the treatment presented here is Boudol and Castellani’s concurrency relation for CCS [1989].




Fig. 5: Conventional diamond property for $t \smile t'$

The essence of $t \smile t'$ is illustrated in Figure 5. If either execution step is taken, the other remains valid, and moreover once both are taken, one ends up in (essentially) the same state, regardless of which step is taken first. However, concurrent transitions are not completely independent: the location and indeed the nature of the redex acted on by one transition may change as a consequence of the earlier transition being taken. This idea is captured by the *residual* t'/t (“ t' after t ”), which specifies how t' must be transformed to operate on $\text{tgt}(t)$ (sometimes called *pseudocommutation* [Angiuli et al. 2014]).

The requirement that t'/t and t/t' are *cofinal* – have the same target state – is straightforward when the transitions preserve the free variables of a term. This is trivially the case in CCS since there are no binders, and is also true of the λ -calculus, where reductions are usually defined on closed terms. In the late-style π -calculus that we consider here, there are transition rules that “open” a process with respect to a name, with the action on the transition representing the upwards propagation of the binder through the process term. In this setting the notion of cofinality is non-trivial; while de Bruijn indices make this subtlety more explicit, we note that the reordering of binders complicates things even in the named setting. We discuss this, with examples, in this section. Permutation of concurrent transitions induces a congruence on traces called *causal equivalence*, which we turn to in § 4.

3.1. Concurrent transitions

In our setting, a transition $t : P \xrightarrow{a} R$ is a *proof* that locates a redex in P , witnessing the fact that $(P, a, R) \in \longrightarrow$. The concurrency relation \smile relates two such proofs; it is defined as the symmetric closure of the relation defined inductively by the rules in Figure 6. The figure makes use of the compact notation for transitions introduced in

§ 2.2. As before, trivial variants of the rules are omitted for clarity. For the transition constructors of the form $\cdot^a | Q$ and $v^a \cdot$ that come in bound and non-bound variants, we abuse notation a little and write a single \smile rule quantified over a to mean that there are two separate (but otherwise identical) cases.

$t \smile t'$

$\frac{}{P ^a u \smile t^{a'} Q}$	$\frac{t \smile t'}{t^a Q \smile t' _y u}$	$\frac{t \smile t'}{t^a Q \smile t' _y u}$	$\frac{u \smile u'}{P ^a u \smile t _y u'}$
$\frac{u \smile u'}{P ^a u \smile t _y u'}$	$\frac{t \smile t'}{t^a Q \smile t' _v u}$	$\frac{t \smile t'}{t^a Q \smile t' _v u}$	$\frac{u \smile u'}{P ^a u \smile t _v u'}$
$\frac{u \smile u'}{P ^a u \smile t _v u'}$	$\frac{t \smile t'}{t + Q \smile t' + Q}$	$\frac{u \smile u'}{P + u \smile P + u'}$	$\frac{t \smile t'}{P ^a t \smile P ^{a'} t'}$
$\frac{t \smile t'}{t^a Q \smile t^{a'} Q}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _z u'}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _z u'}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _z u'}$
$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _v u'}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _v u'}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _v u'}$	$\frac{t \smile t' \quad u \smile u'}{t _y u \smile t' _v u'}$
$\frac{t \smile t' \quad u \smile u'}{t _v u \smile t' _v u'}$	$\frac{t \smile t' \quad u \smile u'}{t _v u \smile t' _v u'}$	$\frac{t \smile t' \quad u \smile u'}{t _v u \smile t' _v u'}$	$\frac{t \smile t'}{\nabla t \smile \nabla t'}$
$\frac{t \smile t'}{\nabla t \smile \nabla t'}$	$\frac{t \smile t'}{v^a t \smile v^{a'} t'}$	$\frac{t \smile t'}{!t \smile !t'}$	

Fig. 6: Concurrent transitions

Intuitively, transitions are concurrent when they pick out non-overlapping redexes. The only axiom, $P |^a u \smile t^{a'} | Q$, says that two transitions t and u are concurrent if they take place on opposite sides of a parallel composition. The remaining rules propagate concurrent sub-transitions up through restriction, choice, parallel composition, and replication. There are no rules allowing us to conclude that a transition which takes the left branch of a choice is concurrent with a transition which takes the right branch of the same choice; choices are mutually exclusive. Likewise, there are no rules allowing us to conclude that an input or output transition is concurrent with any other transition. Since t and t' are coinital, if one of them picks out a prefix then the other necessarily picks out the same prefix, and so they are equal and thus not concurrent.

The $t |_y u \smile t' |_z u'$ rule says that a synchronisation is concurrent with another, as long as the two input transitions t and t' are concurrent on the left branch of the parallel composition, and the two output transitions u and u' are concurrent on the right. The $t |_y u \smile t' |_z u'$ variant is similar, but permits concurrent input and output transitions

on the left, with their respective synchronisation partners concurrent on the right. The $t \mid_y u \smile t' \mid_v u'$ rule and variants are analogous, but permit a plain synchronisation to be concurrent with a v -synchronisation. The main result of this section is that the concurrency relation captured by \smile is sound up to a suitable notion of cofinality.

Example 1 (Concurrent transitions). Consider the π -calculus process (using conventional named syntax) $(\nu y) \bar{x}(y).P \mid \bar{z}(y).Q$. This can take two transitions, the first one sending y on x , resulting in $P \mid \bar{z}(y).Q$, and the second one sending y on channel z , resulting in $\bar{x}(y).P \mid Q$. Notice that y becomes free in both processes.

In de Bruijn notation, this process is written $\nu(\bar{x} + \bar{1}\langle 0 \rangle).P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q$. It can take two transitions, each resulting in an extrusion of the ν -binder; call these t and t' . The transition t initiates the extrusion \bar{x} on the left branch of the parallel composition:

$$\bar{v}. \frac{\bar{x} + \bar{1}\langle 0 \rangle \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \quad \frac{\Gamma + 1 \vdash \bar{x} + \bar{1}\langle 0 \rangle.P \quad \bar{x} + \bar{1}\langle 0 \rangle \rightarrow \Gamma + 1 \vdash P}{\Gamma + 1 \vdash \bar{x} + \bar{1}\langle 0 \rangle.P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{x} + \bar{1}\langle 0 \rangle} \Gamma + 1 \vdash P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q}}{\Gamma \vdash \nu(\bar{x} + \bar{1}\langle 0 \rangle).P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{x}} \Gamma + 1 \vdash P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q}$$

The transition t' initiates an extrusion \bar{z} of the same binder on the right branch of the parallel composition:

$$\bar{v}. \frac{\bar{x} + \bar{1}\langle 0 \rangle.Q \mid \bar{z} + \bar{1}\langle 0 \rangle. \quad \frac{\Gamma + 1 \vdash \bar{z} + \bar{1}\langle 0 \rangle.Q \quad \bar{z} + \bar{1}\langle 0 \rangle \rightarrow \Gamma + 1 \vdash Q}{\Gamma + 1 \vdash \bar{x} + \bar{1}\langle 0 \rangle.P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{z} + \bar{1}\langle 0 \rangle} \Gamma + 1 \vdash \bar{x} + \bar{1}\langle 0 \rangle.P \mid Q}}{\Gamma \vdash \nu(\bar{x} + \bar{1}\langle 0 \rangle).P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{z}} \Gamma + 1 \vdash \bar{x} + \bar{1}\langle 0 \rangle.P \mid Q}$$

Since the two transitions are coinitial and arise on opposite sides of a parallel composition, we can conclude $t \smile t'$ using the rules in Figure 6. Here is the proof, writing the derivations t and t' above using the compact notation for transitions:

$$\frac{\bar{x} + \bar{1}\langle 0 \rangle.P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{x} + \bar{1}\langle 0 \rangle} \bar{x} + \bar{1}\langle 0 \rangle.P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q}{\bar{v}(\bar{x} + \bar{1}\langle 0 \rangle).P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q \xrightarrow{\bar{x}} \bar{v}(\bar{x} + \bar{1}\langle 0 \rangle).P \mid \bar{z} + \bar{1}\langle 0 \rangle.Q} \smile$$

■

3.2. Residuals of concurrent transitions

If two transitions are concurrent then their respective *residuals* provide a canonical way of merging or reconciling them.

Definition 1 (Residual t/t'). For any $t \smile t'$, the *residual* of t after t' , written t/t' , is defined by the equations in Figure 7.

t/t'

$$\begin{aligned}
(P \mid^a u)/(t \mid^c Q) &= \mathbf{tgt}(t) \mid^a u \\
(P \mid^a u)/(t \mid^b Q) &= \mathbf{tgt}(t) \mid^a \mathbf{push}^* u \\
(t \mid^a Q)/(P \mid^c u) &= t \mid^a \mathbf{tgt}(u) \\
(t \mid^a Q)/(P \mid^b u) &= \mathbf{push}^* t \mid^a \mathbf{tgt}(u) \\
(t \mid^a Q)/(t' \mid_y u) &= (\mathbf{pop} \ y)^*(t/t') \mid^a \mathbf{tgt}(u) \\
(P \mid^a u)/(t \mid_y u') &= (\mathbf{pop} \ y)^* \mathbf{tgt}(t) \mid^a u/u' \\
(t \mid_y u)/(t' \mid^b Q) &= t/t' \mid_y \mathbf{push}^* u \\
(t \mid_y u)/(t' \mid^c Q) &= t/t' \mid_y u \\
(t \mid_y u)/(P \mid^b u') &= \mathbf{push}^* t \mid_y u/u' \\
(t \mid_y u)/(P \mid^c u') &= t \mid_y u/u' \\
(t \mid^x Q)/(t' \mid_v u) &= v^x(t/t' \mid^{x+1} \mathbf{tgt}(u)) \\
(t \mid^x Q)/(t' \mid_v u) &= \overline{v}(t/t' \mid^{x+1(0)} \mathbf{tgt}(u)) \\
(t \mid^c Q)/(t' \mid_v u) &= v^c(t/t' \mid^{\mathbf{push}^* c} \mathbf{tgt}(u)) \\
(P \mid^x u)/(t \mid_v u') &= v^x(\mathbf{tgt}(t) \mid^{x+1} u/u') \\
(P \mid^x u)/(t \mid_v u') &= \overline{v}(\mathbf{tgt}(t) \mid^{x+1(0)} u/u') \\
(P \mid^c u)/(t \mid_v u') &= v^c(\mathbf{tgt}(t) \mid^{\mathbf{push}^* c} u/u') \\
(t \mid_v u)/(t' \mid^b Q) &= t/t' \mid_v \mathbf{push}^* u \\
(t \mid_v u)/(t' \mid^c Q) &= t/t' \mid_v u \\
(t \mid_v u)/(P \mid^x u') &= \mathbf{push}^* t \mid_v u/u' \\
(t \mid_v u)/(P \mid^x u') &= \mathbf{push}^* t \mid_v u/u' \\
(t \mid_v u)/(P \mid^c u') &= t \mid_v u/u' \\
(t + Q)/(t' + Q) &= t/t' \\
(P \mid^x u)/(P \mid^b u') &= \mathbf{push}^* P \mid^x u/u' \\
(P \mid^b u)/(P \mid^x u') &= \mathbf{push}^* P \mid^b u/u' \\
(P \mid^x u)/(P \mid^y u') &= \mathbf{push}^* P \mid^{x+y(0)} u/u' \\
(P \mid^c u)/(P \mid^b u') &= \mathbf{push}^* P \mid^c u/u' \\
(P \mid^a u)/(P \mid^c u') &= P \mid^a u/u' \\
(t \mid^x Q)/(t' \mid^b Q) &= t/t' \mid^x \mathbf{push}^* Q \\
(t \mid^b Q)/(t' \mid^x Q) &= t/t' \mid^b \mathbf{push}^* Q \\
(t \mid^x Q)/(t' \mid^y Q) &= t/t' \mid^{x+y(0)} \mathbf{push}^* Q \\
(t \mid^c Q)/(t' \mid^b Q) &= t/t' \mid^c \mathbf{push}^* Q \\
(t \mid^a Q)/(t' \mid^c Q) &= t/t' \mid^a Q \\
(t \mid_y u)/(t' \mid_z u') &= (\mathbf{pop} \ z)^*(t/t') \mid_y u/u' \\
(t \mid_y u)/(t' \mid_v u') &= v^x(t/t' \mid_y u/u') \\
(t \mid_v u)/(t' \mid_z u') &= (\mathbf{pop} \ z)^*(t/t') \mid_v u/u' \\
(t \mid_v u)/(t' \mid_v u') &= v^x(t/t' \mid_v u/u') \\
(\overline{v}t)/(\overline{v}t') &= t/t' \\
(\overline{v}t)/(\overline{v}^b t') &= \overline{v} \mathbf{swap}^*(t/t') \\
(\overline{v}t)/(\overline{v}^c t') &= \overline{v} \ t/t' \\
(\overline{v}^b t)/(\overline{v}t') &= t/t' \\
(\overline{v}^c t)/(\overline{v}t') &= t/t' \\
(\overline{v}^b t)/(\overline{v}^b t') &= v \ t/t' \\
(\overline{v}^c t)/(\overline{v}^b t') &= v^c \mathbf{swap}^*(t/t') \\
(\overline{v}^b t)/(\overline{v}^c t') &= v^b \ t/t' \\
(\overline{v}^c t)/(\overline{v}^c t') &= v^c \ t/t' \\
(!t)/(!t') &= t/t'
\end{aligned}$$

Fig. 7: Residual of t after t' , omitting $\cdot_y \mid \cdot$ and $\cdot_v \mid \cdot$ cases

The above definition is a total and terminating function on concurrent transitions; in Agda, this is verified by the typechecker. Syntactically, the operator \cdot/\cdot has higher precedence than any transition constructor. The definition makes use of the renaming lemmas in § 2.1.2 and the fact that the transition system is closed under renamings (Lemma 9).

While the definition is rather technical, the idea is quite simple: the residual says how to update one transition to take into account the fact that the other has taken place, for example by adjusting the path to the redex, or applying an appropriate renaming. Several examples are included in the sections which follow. Example 2 below gives the basic idea, and § 3.3, which explains the notion of cofinality, shows how these “residual redexes” are obtained in more complicated cases.

Example 2 (Residuals of concurrent transitions). First recall the named process in Example 1 above, that is, $(vy) \overline{x}(y).P \mid \overline{z}(y).Q$. Both of the transitions it can perform are bound transitions, extruding y , which is no longer bound in the resulting process. After the first transition, the second can be performed as a free send of y along z and vice versa, and in both cases we obtain the process $P \mid Q$, again containing y free.

These observations are reflected in the de Bruijn representation. Since t and t' are concurrent there should exist residual transitions denoted t'/t and t/t' , which are cofinal, allowing us to complete the square

$$\begin{array}{ccc}
 & \Gamma + 1 \vdash P \mid \overline{z+1}\langle 0 \rangle . Q & \\
 t \nearrow & & \searrow t'/t \\
 \Gamma \vdash \nu(\overline{x+1}\langle 0 \rangle . P \mid \overline{z+1}\langle 0 \rangle . Q) & & \Gamma + \Delta \vdash R \\
 t' \searrow & & \nearrow t/t' \\
 & \Gamma + 1 \vdash \overline{x+1}\langle 0 \rangle . P \mid Q &
 \end{array}$$

for some $\Delta \in \{1, 2\}$ and some process R . In the upper state $\text{tgt}(t)$, the only candidate for t'/t is the output prefix $\overline{z+1}\langle 0 \rangle$. However, the ν -binder to which index 0 refers no longer appears in $\text{tgt}(t)$. Rather, that binder is propagating and index 0 is free, reflected by $\text{tgt}(t)$ being in context $\Gamma + 1$. When the output transition is taken, $\overline{z+1}\langle 0 \rangle$ therefore simply propagates as a non-bound action, rather than causing a further extrusion:

$$\begin{array}{c}
 P \mid \overline{z+1}\langle 0 \rangle . \frac{\Gamma + 1 \vdash \overline{z+1}\langle 0 \rangle . Q \xrightarrow{\overline{z+1}\langle 0 \rangle} \Gamma + 1 \vdash Q}{\Gamma + 1 \vdash P \mid \overline{z+1}\langle 0 \rangle . Q \xrightarrow{\overline{z+1}\langle 0 \rangle} \Gamma + 1 \vdash P \mid Q}
 \end{array}$$

From the lower state $\text{tgt}(t')$ the only candidate for t/t' is the output prefix $\overline{x+1}\langle 0 \rangle$, and similar reasoning applies. Thus for t/t' we have

$$\begin{array}{c}
 \frac{\overline{x+1}\langle 0 \rangle \mid Q \frac{\Gamma + 1 \vdash \overline{x+1}\langle 0 \rangle . P \xrightarrow{\overline{x+1}\langle 0 \rangle} \Gamma + 1 \vdash P}{\Gamma + 1 \vdash \overline{x+1}\langle 0 \rangle . P \mid Q \xrightarrow{\overline{x+1}\langle 0 \rangle} \Gamma + 1 \vdash P \mid Q}}{\Gamma + 1 \vdash \overline{x+1}\langle 0 \rangle . P \mid Q}
 \end{array}$$

and therefore $\Delta = 1$ and $R = P \mid Q$. In summary when concurrent t and t' extrude the same binder, their respective residuals are plain outputs, not bound outputs, because a given binder can only be extruded once.

To relate this example to the defining equations of \cdot/\cdot we use the compact presentations of t and t' from the end of Example 1. It is then easy to see that the rules in Figure 7 indeed compute (in compact form) the derivation above for t/t' :

$$\begin{aligned}
 & \overline{\nu(x+1)\langle 0 \rangle} . P \mid \overline{z+1}\langle 0 \rangle . Q / \overline{\nu(x+1)\langle 0 \rangle} . P \mid \overline{z+1}\langle 0 \rangle . Q \\
 = & \overline{(x+1)\langle 0 \rangle} . P \mid \overline{z+1}\langle 0 \rangle . Q / \overline{(x+1)\langle 0 \rangle} . P \mid \overline{z+1}\langle 0 \rangle . Q \\
 = & \overline{x+1}\langle 0 \rangle . P \mid Q
 \end{aligned}$$

and similarly for t'/t . ■

Example 2 illustrated the basic idea of residuation, focusing on the specific case where the residuals of transitions with bound actions have actions that are not bound, a subtlety of residuation particular to π -calculus first noted by Cristescu et al. [2013]. To capture this and other aspects of residuation, it is useful to define a datatype of *concurrent actions* $a \smile a'$ and an associated notion of residual action a/a' and use these to index concurrent transitions and their residuals.

We define both of these using the diagrams in Figure 8 below. The datatype of concurrent actions, ranged over by \hat{a} , has five constructors, one for each diagram; the arrows

diverging on the left represent the concurrent actions a and a' , and the arrows converging on the right define the corresponding residuals a'/a and a/a' . Beneath each diagram is the *braiding* relation \bowtie_a which constitutes the notion of cofinality induced by that form of concurrent action.




Fig. 8: Concurrent actions $\dot{a} : a \smile a'$, residuals a'/a and a/a' , and braiding relation \bowtie_a

Diagrams (i) and (ii) capture the general pattern when at most one of the actions is bound. In (ii), image of an action in a bound action is the original action shifted under a binder; in both cases cofinality is simply equality. Diagram (iii) is the general pattern when both actions are bound: in this case the target states P and P' are related by a “free braid” $P \bowtie P'$ in the form of the permutation **swap** which renames 0 to 1 and 1 to 0, reflecting the transposition of the two binders. Free braids are illustrated in some detail in Examples 3 and 4 below.

Diagram (iv) and (v) are specific to name extrusion. Diagram (iv) is an exception to (iii) where the two bound actions happen to be extrusions of the same binder, as in Example 1; the $\overline{\nu}t \smile \overline{\nu}t'$ rule is the only concurrency rule that generates concurrent actions of this form. Diagram (v) is an exception to (i) where the two non-bound actions happen to be ν -synchronisations of distinct binders. In this case the residual actions will also be ν -synchronisations (as suggested by the informal τ_ν notation) and the target states are related by a “bound braid” $P \bowtie P'$, essentially a free braid which has been “closed” by a pair of ν -binders, representing the transposition of those binders. The four variants of the $t \mid_\nu u \smile t' \mid_\nu u'$ rule generate concurrent actions of this form whenever the extruding binders are distinct. Bound braids are illustrated in Example 5 below.

Free and bound braids are now defined more formally.

Definition 2 (Free braid). For any processes $\Gamma + 2 \vdash P, R$ define the symmetric relation $P \bowtie R$ as follows. The context Γ is left implicit.

$$P \bowtie R \iff P = \text{swap}^* R$$

Symmetry of \bowtie is immediate from the involutivity of **swap**. Note that \bowtie is not irreflexive, since $\text{swap}^* P = P$ iff indices 0 and 1 are both unused in P .

Definition 3 (Bound braid). For any processes $\Gamma \vdash P, R$ inductively define the symmetric relation $P \bowtie R$ using the rules in Figure 9. Again the context Γ is left implicit.

$$\begin{array}{c}
\boxed{P \times R} \\
\\
\text{vv-swap}_P \frac{P \times R}{\text{vv}P \times \text{vv}R} \quad \cdot + Q \frac{P \times R}{P + Q \times R + Q} \quad P + \cdot \frac{Q \times S}{P + Q \times P + S} \\
\\
\cdot | Q \frac{P \times R}{P | Q \times R | Q} \quad P | \cdot \frac{Q \times S}{P | Q \times P | S} \quad \nu \cdot \frac{P \times R}{\nu P \times \nu R} \quad ! \cdot \frac{P \times R}{!P \times !R}
\end{array}$$

Fig. 9: Bound braid $P \times R$

Note that the vv-swap_P rule requires P and R to be related by a *free* braid ($P \times R$) which it then closes with a pair of ν -binders. By contrast the $\nu \cdot$ rule simply propagates a bound braid ($P \times R$) through a ν -binder.

We adopt a compact term-like notation for \times proofs similar to the convention introduced earlier for transitions. As before, rule names are shown to the left of each rule, in blue. The symmetry of \times follows easily from the symmetry of \times ; moreover \times is also not irreflexive, because \times is not irreflexive. Meta-variables ϕ and ψ range over bound braids; $\text{src}(\phi)$ and $\text{tgt}(\phi)$ denote P and R for any $\phi : P \times R$. Bound braids are “unobservable” in the sense that two processes related by a bound braid are strongly bisimilar. Indeed $\text{vv}(\text{swap}^*P) \cong \text{vv}P$ is simply the de Bruijn counterpart of the familiar congruence $(\nu xy) P \cong (\nu yx) P$. However in our constructive setting – at least in the absence of non-trivial techniques or extensions to type theory – the usual refrain “work up to structural congruence!” is of little help; representing such congruences would still require explicit witnesses at least as complex as bound braids.

Concurrent actions and action residuals give transition residuals a more precise type (omitted for simplicity from the definitions of $t \sim t'$ and t/t'), making them somewhat easier to formally define. But more important here is how they determine the appropriate notion of cofinality relating $\text{tgt}(t'/t)$ and $\text{tgt}(t/t')$, namely the braiding relation $\mathbf{x}_{a,a'}$ specified beneath each diagram in Figure 8. A braiding relation is a singleton type, whose unique inhabitant precisely captures precisely the “rewiring” effect of reordering transitions that involve binders.

Definition 4 (Braiding). For any context Γ , any $a, a' \in \text{Action } \Gamma$ and any $\hat{a} : a \sim a'$, define the following symmetric relation $\mathbf{x}_{\hat{a}}$ over processes in Γ' , where Γ' is the target context of \hat{a} .

$$\mathbf{x}_{\hat{a}} \stackrel{\text{def}}{=} \text{one of } \times, \times \text{ or } = \text{ as defined in Figure 8}$$

Our key soundness result is that the targets of the residuals of concurrent transitions $t \sim t'$ with actions $\hat{a} : a \sim a'$ are cofinal in the sense of being related by $\mathbf{x}_{\hat{a}}$. We need first that bound braids are closed under renamings, which we capture as a notion of residuation ϕ/ρ . The other residual ρ/ϕ is always ρ .

Lemma 10. For any $\Gamma \vdash P$, suppose $\phi : P \times Q$ and $\rho : \Gamma \longrightarrow \Delta$. Then there exists a bound braid $\phi/\rho : \rho^*P \times \rho^*Q$.

$$\begin{array}{ccc} P & \xrightarrow{\phi} & Q \\ \rho^* \downarrow & & \downarrow \rho^* \\ \rho^*P & \xrightarrow{\phi/\rho} & \rho^*Q \end{array}$$

Theorem 1 (Cofinality of residuals).

$$\begin{array}{ccc} & \Gamma' \vdash R & \xrightarrow{t'/t} \Gamma + \Delta \vdash Q \\ \Gamma \vdash P & \xrightarrow{t} & \\ & \Gamma'' \vdash R' & \xrightarrow{t/t'} \Gamma + \Delta \vdash Q' \end{array} \quad \begin{array}{c} \text{---} \\ \gamma_{t,t'} \end{array}$$

Suppose $t \smile t'$ with actions $\dot{a} : a \smile a'$. Then there exists a unique $\gamma_{t,t'} : \text{tgt}(t'/t) \times_{\dot{a}} \text{tgt}(t/t')$.

We omit the t, t' subscripts when the particular concurrent actions are immaterial.

There is no analogous result to show that the definition of concurrency is complete: that it includes every pair of coinital transitions for which a cofinal notion of residuation might be defined. It is not entirely clear what form such a theorem might take, nor are we aware of any such theorem in the literature. Choice in particular is potentially problematic. Although by our (and Boudol and Castellani's) definition of \smile coinital choices are never concurrent, in the following we have distinct coinital choices with “obvious” residuals, which are indeed cofinal:

$$\begin{array}{ccc} & \underline{x}.P & \\ \cdot + \underline{x}.P & \nearrow & \searrow \underline{x}.P \\ (\underline{x}.P) + \underline{x}.P & & P \\ (\underline{x}.P) + \cdot & \searrow & \nearrow \underline{x}.P \\ & \underline{x}.P & \end{array}$$

One avenue for justifying this (and other) choices about the transition concurrency relation for π -calculus might be to prove results about the equivalence of “proved transition” semantics and event structure semantics, analogous to the results of Boudol and Castellani [1991] for CCS. We leave exploring canonical notions of concurrency to future work.

3.3. Examples of braiding

Example 3 (Free braid).

Free braids arise when there are concurrent bound actions. For example the push injections used in the propagation rules $\cdot \stackrel{\lambda}{|} Q$ and $P \stackrel{\lambda}{|} \cdot$ open the process term with

respect to index 0; if two of these happen consecutively, the order in which they happen determines the roles of indices 0 and 1 in the final process.

Concurrent name extrusions are analogous. In the process term $\nu\nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q$ there are two binders that can be extruded; call the outer one ν_1 and the inner one ν_2 . The output on the left extrudes ν_2 , and the output on the right extrudes ν_1 . Let t be the transition that extrudes ν_2 :

$$\begin{array}{c} \frac{\frac{\frac{\overline{x+2}\langle 0 \rangle \mid \overline{z+2}\langle 1 \rangle.Q}{\overline{x+2}\langle 0 \rangle} \quad \frac{\Gamma + 2 \vdash \overline{x+2}\langle 0 \rangle.P}{\Gamma + 2 \vdash P}}{\Gamma + 2 \vdash \overline{x+2}\langle 0 \rangle.P \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{x+2}\langle 0 \rangle} \Gamma + 2 \vdash P \mid \overline{z+2}\langle 1 \rangle.Q}{\overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{x+2}\langle 0 \rangle} \Gamma + 2 \vdash P \mid \overline{z+2}\langle 1 \rangle.Q \\ \frac{\Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q}{\Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{x+1}} \Gamma + 2 \vdash P \mid \overline{z+2}\langle 1 \rangle.Q}{\Gamma \vdash \nu\nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{x}} \Gamma + 1 \vdash \nu(\text{swap}^*P \mid \overline{z+2}\langle 0 \rangle.\text{swap}^*Q) \end{array}$$

Here ν_2 is extruded as the bound output $\overline{x+1}$, propagating through the outer binder ν_1 as \overline{x} . In $\text{tgt}(t)$, index 0 refers to the extruding ν_2 ; the binder remaining in the process term is ν_1 . The key detail here is that the rule $\nu^{\overline{x}}$ moves ν_1 past ν_2 , explaining the use of **swap** in $\text{tgt}(t)$: whenever a propagating binder moves past a static binder, a **swap** must be applied under the static binder to preserve the local meaning of indices 0 and 1 (§ 2.2 above). This is also why $\overline{z+2}\langle 1 \rangle$ in $\text{src}(t)$ becomes $\overline{z+2}\langle 0 \rangle$ in $\text{tgt}(t)$.

Now let t' be the transition that extrudes ν_1 :

$$\begin{array}{c} \frac{\frac{\frac{(\overline{x+2}\langle 0 \rangle).P \mid \overline{z+2}\langle 1 \rangle.Q}{\overline{x+2}\langle 0 \rangle} \quad \frac{\Gamma + 2 \vdash \overline{z+2}\langle 1 \rangle.Q}{\Gamma + 2 \vdash Q}}{\Gamma + 2 \vdash (\overline{x+2}\langle 0 \rangle).P \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{z+2}\langle 1 \rangle} \Gamma + 2 \vdash Q}{\overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{z+2}\langle 1 \rangle} \Gamma + 2 \vdash (\overline{x+2}\langle 0 \rangle).P \mid Q \\ \frac{\Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q}{\Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{z+1}\langle 0 \rangle} \Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid Q}{\Gamma \vdash \nu\nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q} \xrightarrow{\overline{z}} \Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid Q \end{array}$$

In this case, the output $\overline{x+2}\langle 1 \rangle$ propagates through the inner binder ν_2 as $\overline{z+1}\langle 0 \rangle$, and then becomes the extrusion \overline{z} of the outer binder ν_1 . In $\text{tgt}(t')$, index 0 thus refers to the extruding ν_1 , and the binder that remains in the process term is ν_2 . The key detail here is that there is no **swap** in $\text{tgt}(t')$ because this time the relative positions of ν_1 and ν_2 are unchanged: the extruding ν_1 is still the outer of the two binders.

The proof that t and t' are concurrent is straightforward because the outputs occur under opposite sides of a parallel composition. The notion of cofinality, however, is complicated by the use of indices to refer to ν_1 and ν_2 . Naively, our expectation would be to derive t'/t and t/t' that complete the square

$$\begin{array}{ccc} & \Gamma + 1 \vdash \nu(\text{swap}^*P \mid \overline{z+2}\langle 0 \rangle.\text{swap}^*Q) & \\ & \nearrow t \quad \quad \quad \searrow t'/t & \\ \Gamma \vdash \nu\nu((\overline{x+2}\langle 0 \rangle).P) \mid \overline{z+2}\langle 1 \rangle.Q & & \Gamma + \Delta \vdash R \\ & \searrow t' \quad \quad \quad \nearrow t/t' & \\ & \Gamma + 1 \vdash \nu((\overline{x+2}\langle 0 \rangle).P) \mid Q & \end{array}$$

for some $\Delta \in \{1, 2\}$ and some process R . However, the only candidate for t'/t is to select

the output redex on the right, which becomes an extrusion of the remaining binder, in this case v_1 :

$$\overline{v}. \frac{\text{swap}^* P \mid \overline{z+2(0)}. \frac{\Gamma + 2 \vdash \overline{z+2(0)}. \text{swap}^* Q \xrightarrow{\overline{z+2(0)}} \Gamma + 2 \vdash \text{swap}^* Q}{\Gamma + 2 \vdash \text{swap}^* P \mid \overline{z+2(0)}. \text{swap}^* Q \xrightarrow{\overline{z+2(0)}} \Gamma + 2 \vdash \text{swap}^* P \mid \text{swap}^* Q}}{\Gamma + 1 \vdash v(\text{swap}^* P \mid \overline{z+2(0)}. \text{swap}^* Q) \xrightarrow{\overline{z+2}} \Gamma + 2 \vdash \text{swap}^* P \mid \text{swap}^* Q}$$

leaving indices 0, 1 referring to v_1, v_2 respectively in $\text{tgt}(t'/t)$. Equally, the only candidate for the other residual t/t' is to select the output redex on the left, which also becomes an extrusion of the remaining binder, in this case v_2 :

$$\overline{v}. \frac{\overline{x+2(0)} \mid Q \quad \frac{\Gamma + 2 \vdash \overline{x+2(0)}. P \xrightarrow{\overline{x+2(0)}} \Gamma + 2 \vdash P}{\Gamma + 2 \vdash (\overline{x+2(0)}. P) \mid Q \xrightarrow{\overline{x+2(0)}} \Gamma + 2 \vdash P \mid Q}}{\Gamma + 1 \vdash v(\overline{x+2(0)}. P) \mid Q \xrightarrow{\overline{x+2}} \Gamma + 2 \vdash P \mid Q}$$

leaving indices 0, 1 in $\text{tgt}(t/t')$ referring to v_2, v_1 rather than v_1, v_2 . So instead of the expected square, we have the pentagon

$$\begin{array}{ccc} \Gamma + 1 \vdash v(\text{swap}^* P \mid \overline{z+2(0)}. \text{swap}^* Q) & \xrightarrow{t'/t} & \Gamma + 2 \vdash \text{swap}^* P \mid \text{swap}^* Q \\ \nearrow t & & \downarrow \text{swap}^* \\ \Gamma \vdash v v(\overline{x+2(0)}. P) \mid \overline{z+2(1)}. Q & & \\ \searrow t' & & \\ \Gamma + 1 \vdash v(\overline{x+2(0)}. P) \mid Q & \xrightarrow{t/t'} & \Gamma + 2 \vdash P \mid Q \end{array}$$

with a swap path between $\text{tgt}(t'/t)$ and $\text{tgt}(t/t')$ reflecting the reordering of the propagating binders v_1 and v_2 . \blacksquare

Example 4 (Propagating free braid).

Free braids are preserved by enclosing transitions as long as the residual actions of those transitions are bound. In particular, if a free braid propagates through a v -binder it remains a free braid. Suppose $t \smile t'$ where the residual actions are both bound, so that $\text{tgt}(t'/t) \times \text{tgt}(t/t')$:

$$\begin{array}{ccc} \Gamma + 2 \vdash R & \xrightarrow{(t'/t)^{\overline{z+2}}} & \Gamma + 3 \vdash P' \\ \nearrow t^{\overline{x+1}} & & \downarrow \text{swap}^* \\ \Gamma + 1 \vdash P & & \\ \searrow t'^{\overline{z+1}} & & \\ \Gamma + 2 \vdash R' & \xrightarrow{(t/t')^{\overline{x+2}}} & \Gamma + 3 \vdash \text{swap}^* P' \end{array}$$

Since both $\overline{x+1}$ and $\overline{z+1}$ are of the form $\text{push}^* b$, we can use the $\overline{v}^b.$ rule to form the

composite transitions $\nu^{\underline{x}}t$ and $\nu^{\underline{z}}t'$ which propagate the input actions of t and t' actions through a ν -binder:

$$\nu^{\underline{x}}. \frac{t \frac{\vdots}{\Gamma + 1 \vdash P \xrightarrow{\underline{x}+1} \Gamma + 2 \vdash R}}{\Gamma \vdash \nu P \xrightarrow{\underline{x}} \Gamma + 1 \vdash \nu(\text{swap}^* R)} \quad \nu^{\underline{z}}. \frac{t' \frac{\vdots}{\Gamma + 1 \vdash P \xrightarrow{\underline{z}+1} \Gamma + 2 \vdash R'}}{\Gamma \vdash \nu P \xrightarrow{\underline{z}} \Gamma + 1 \vdash \nu(\text{swap}^* R')}$$

Since $t \smile t'$ we can conclude $\nu^{\underline{x}}t \smile \nu^{\underline{z}}t'$ by the rules in Figure 7 and compute the following composite residual $(\nu^{\underline{z}}t')/\nu^{\underline{x}}t$:

$$\nu^{\underline{z}+1}. \frac{\text{swap}^*. \frac{t'/t \frac{\vdots}{\Gamma + 2 \vdash R \xrightarrow{\underline{z}+2} \Gamma + 3 \vdash P'}}{\Gamma + 2 \vdash \text{swap}^* R \xrightarrow{\underline{z}+2} \Gamma + 3 \vdash (\text{swap} + 1)^* P'}}{\Gamma + 1 \vdash \nu(\text{swap}^* R) \xrightarrow{\underline{z}+1} \Gamma + 2 \vdash \nu(\text{swap}^*(\text{swap} + 1)^* P')}$$

noting that $\text{swap}^*(\underline{z} + 2) = \underline{z} + 2$ by Lemma 8. The other residual $(\nu^{\underline{x}}t)/\nu^{\underline{z}}t'$ is similar but has an extra swap inherited from $\text{tgt}(t/t')$:

$$\nu^{\underline{x}+1}. \frac{\text{swap}^*. \frac{t/t' \frac{\vdots}{\Gamma + 2 \vdash R' \xrightarrow{\underline{x}+2} \Gamma + 3 \vdash \text{swap}^* P'}}{\Gamma + 2 \vdash \text{swap}^* R' \xrightarrow{\underline{x}+2} \Gamma + 3 \vdash (\text{swap} + 1)^* \text{swap}^* P'}}{\Gamma + 1 \vdash \nu(\text{swap}^* R') \xrightarrow{\underline{x}+1} \Gamma + 2 \vdash \nu(\text{swap}^*(\text{swap} + 1)^* \text{swap}^* P')}$$

Nevertheless, the target states of the composite residuals are still equated by swap , consistent with the fact that the residual actions still bound.

$$\begin{array}{ccc} \Gamma + 1 \vdash \nu(\text{swap}^* R) & \xrightarrow{\nu^{\underline{z}+1}(\text{swap}^* t'/t)} & \Gamma + 2 \vdash \nu(\text{swap}^*(\text{swap} + 1)^* P') \\ \nearrow \nu^{\underline{x}}t & & \downarrow \text{swap}^* \\ \Gamma \vdash \nu P & & \Gamma + 2 \vdash \nu((\text{swap} + 1)^* \text{swap}^*(\text{swap} + 1)^* P') \\ \searrow \nu^{\underline{z}}t' & & \parallel \nu\alpha \\ \Gamma + 1 \vdash \nu(\text{swap}^* R') & \xrightarrow{\nu^{\underline{x}+1}(\text{swap}^* t/t')} & \Gamma + 2 \vdash \nu(\text{swap}^*(\text{swap} + 1)^* \text{swap}^* P') \end{array}$$

Here α is the hexagon equating two ways of transposing indices 0 and 2 (Lemma 3) which $\nu\alpha$ lifts via congruence to an equality between one target and the swap image of the other. Thus \bowtie remains the appropriate notion of cofinality. \blacksquare

Example 5 (Bound braid).

A bound braid arises when concurrent ν -synchronisations have residuals that also ν -synchronise, which requires the underlying extrusions to be distinct binders. The concurrent transitions $t \smile t'$ and $u \smile u'$ below can be composed into concurrent ν -synchronisations that have this property; u has an input \underline{x} matching the bound output \bar{x}

of t , and u' has a bound output \bar{z} matching the input \underline{z} of t' . The extrusions \bar{x} and \bar{z} are clearly of distinct binders since they arise on opposite sides of a parallel composition.

$$\begin{array}{ccc}
 \Gamma \vdash P & \begin{array}{c} \xrightarrow{t^{\bar{x}}} \Gamma + 1 \vdash R \xrightarrow{(t'/t)^{\bar{z}+1}} \Gamma + 2 \vdash P' \\ \xrightarrow{t^{\bar{z}}} \Gamma + 1 \vdash R' \xrightarrow{(t/t')^{\bar{x}+1}} \Gamma + 2 \vdash \text{swap}^* P' \end{array} & \text{swap}^* \\
 & & \\
 \Gamma \vdash Q & \begin{array}{c} \xrightarrow{u^{\bar{x}}} \Gamma + 1 \vdash S \xrightarrow{(u'/u)^{\bar{z}+1}} \Gamma + 2 \vdash Q' \\ \xrightarrow{u^{\bar{z}}} \Gamma + 1 \vdash S' \xrightarrow{(u/u')^{\bar{x}+1}} \Gamma + 2 \vdash \text{swap}^* Q' \end{array} & \text{swap}^*
 \end{array}$$

The composites are the ν -synchronisations $t \mid u : P \mid Q \xrightarrow{\tau} \nu(R \mid S)$ and $t' \mid u' : P' \mid Q' \xrightarrow{\tau} \nu(R' \mid S')$. Moreover since $t \sim t'$ and $u \sim u'$ we can conclude $t \mid u \sim t' \mid u'$ using the rules in Figure 6. The equations in Figure 7 determine the residual $(t' \mid u') / (t \mid u) = \nu^\tau(t'/t \mid u'/u)$, which we write down in full for clarity:

$$\begin{array}{c}
 \vdots \\
 \frac{t'/t \xrightarrow{\bar{z}+1} \Gamma + 1 \vdash S \xrightarrow{\bar{z}+1} \Gamma + 2 \vdash Q' \quad u'/u \xrightarrow{\bar{x}+1} \Gamma + 1 \vdash R \xrightarrow{\bar{x}+1} \Gamma + 2 \vdash P'}{\Gamma + 1 \vdash R \mid S \xrightarrow{\tau} \Gamma + 1 \vdash \nu(P' \mid Q')} \\
 \vdots \\
 \frac{\cdot \mid_\nu \cdot}{\Gamma \vdash \nu(R \mid S) \xrightarrow{\tau} \Gamma \vdash \nu\nu(P' \mid Q')}
 \end{array}$$

The other residual $(t' \mid u') / (t \mid u) = \nu^\tau(t'/t \mid u'/u)$ is similar, except it inherits two extra swap renamings from t/t' and u/u' :

$$\begin{array}{c}
 \vdots \\
 \frac{t/t' \xrightarrow{\bar{x}+1} \Gamma + 1 \vdash S' \xrightarrow{\bar{x}+1} \Gamma + 2 \vdash \text{swap}^* Q' \quad u/u' \xrightarrow{\bar{z}+1} \Gamma + 1 \vdash R' \xrightarrow{\bar{z}+1} \Gamma + 2 \vdash \text{swap}^* P'}{\Gamma + 1 \vdash R' \mid S' \xrightarrow{\tau} \Gamma + 1 \vdash \nu(\text{swap}^* P' \mid \text{swap}^* Q')} \\
 \vdots \\
 \frac{\cdot \mid_\nu \cdot}{\Gamma \vdash \nu(R' \mid S') \xrightarrow{\tau} \Gamma \vdash \nu\nu(\text{swap}^* P' \mid \text{swap}^* Q')}
 \end{array}$$


Thus each residual ν -synchronises, and then propagates through the binder reinserted by the first synchronisation, leaving a double- ν in the final process. The residuals are related by the pentagon

$$\begin{array}{ccc}
 \Gamma \vdash \nu(R \mid S) & \xrightarrow{\nu^\tau(t'/t \mid u'/u)} & \Gamma \vdash \nu\nu(P' \mid Q') \\
 \begin{array}{c} \nearrow t \mid u \\ \searrow t' \mid u' \end{array} & & \downarrow \nu\nu\text{-swap}_{P' \mid Q'} \\
 \Gamma \vdash P \mid Q & & \\
 \begin{array}{c} \nearrow t \mid u \\ \searrow t' \mid u' \end{array} & & \Gamma \vdash \nu(R' \mid S') \xrightarrow{\nu^\tau(t'/t \mid u'/u)} \Gamma \vdash \nu\nu(\text{swap}^* P' \mid \text{swap}^* Q')
 \end{array}$$

where $\nu\nu\text{-swap}_{P' \mid Q'}$ is the bound braid that locates $P' \mid Q' \times \text{swap}^* P' \mid \text{swap}^* Q'$ under the two binders, representing the reordering of the binders. \blacksquare


Example 6 (Braid erasure by synchronisation). A free braid is erased if it is enclosed by a concurrent transition where the notion of cofinality is equality. For example, consider a

variant of Example 5 where the extrusions \bar{x} and \bar{z} occur on the same side of the parallel composition, and represent extrusions of the same binder.



(Using named syntax, the term Q might be of the form $(\nu y) \bar{x}\langle y \rangle.Q_1 \mid \bar{z}\langle y \rangle.Q_2$, as per Example 1 above.)

The residuals u'/u and u/u' are plain outputs, rather than bound outputs. While the composites $t|_{\nu} u \smile t'|_{\nu} u'$ are concurrent ν -synchronisations as before, the residuals of the composites are plain synchronisations, again propagated through the ν -binder reinserted by the preceding step.



Since the residual actions are plain τ actions, cofinality is simply equality. And indeed the substitution $\text{pop } 0$ erases the free braid relating P' and $\text{swap}^* P'$, by mapping indices 0 and 1 both to 0. Here α is the equality $(\text{pop } 0) \circ \text{swap} = \text{pop } 0$ (Lemma 4) and $\nu(\alpha^* P'_{=} \mid Q'_{=})$ uses congruence to lift α to an equivalence on target states, where $P'_{=}$ and $Q'_{=}$ denote the canonical reflexivity proofs of P' and Q' . \blacksquare

This completes our formal treatment of concurrent transitions in π -calculus, including the counterpart of the diamond lemma. In our setting, transitions may open terms with respect to variables, leading to a non-trivial notion of cofinality when such transitions are reordered. Like Boudol and Castellani, we omit a formalisation of Lévy's “cube” property, which extends the notion of concurrency to dimensions greater than two, since it is not required for the formalisation of causal equivalence.

4. Causal equivalence

We now turn to formalising *causal equivalence*, the congruence over sequences of transitions, or *traces*, induced by the concurrency relation for transitions. This is a standard concept from the theory of concurrent alphabets [Mazurkiewicz 1987], but is non-trivial

in our setting because of braidings, which (as we shall see below) both propagate horizontally and compose vertically.

An “atom” of causal equivalence equates $t \cdot t'/t$ and $t' \cdot t/t'$ for concurrent transitions $t \smile t'$, where $t \cdot u$ denotes the composition of t and u . When the associated pentagon is composed horizontally into a larger computation, the continuation must be transported through the braiding $\gamma_{t,t'}$ which relates the target states of t'/t and t/t' . This requires two dimensions of closure, as illustrated in Figure 10. For coinital u and $\gamma_{t,t'}$, the transition u must have an image $u/\gamma_{t,t'}$ in $\gamma_{t,t'}$, and the braiding $\gamma_{t,t'}$ must propagate as $\gamma_{t,t'}/u$:




Fig. 10: Closure of transitions under braidings

The residual $u/\gamma_{t,t'}$ is a version of u which takes into account any braiding that arises from the concurrency of t and t' , whereas $\gamma_{t,t'}/u$ represents the effect of the braiding on the transition u .

For braidings to be preserved by transitions and vice-versa requires two generalisations to the notion of braiding (Definition 4). For free braids, we need the renaming to be of the form $\text{swap} + \Delta$ rather than swap , so that braids can be preserved by subsequent bound actions which further open up the process term. For bound braids, the effect of doing more computation is that the unique pair of binders picked out by a bound braid (Definition 3) may end up being *dropped* (if it occurs on the discarded side of a choice) or *duplicated* (if it occurs under a replication). This requires a more general notion of bound braid closed under reflexivity and parallel composition.

An additional requirement is that braidings compose vertically when causal equivalences are composed via transitivity:




Fig. 11: Sequential composition of concurrent transitions

This diagram represents the causal equivalence

$$t \cdot t'/t \cdot u \cdot u'/u \simeq t' \cdot t/t' \cdot u'/y_{t,t'} \cdot (u/u')/(y_{t,t'}/u')$$

with the targets S and S'' related by the composite braiding $y_{u,u'} \cdot y_{t,t'}/(y_{t,t'}/u')/(u/u')$. It is worth reiterating that while the complexity of tracking free braids is unique to the de Bruijn setting, the implications of bound braids are not, since they arise from transposed binders.

We proceed by defining traces \mathbf{t} (§4.1), and then showing that, suitably generalised, braidings γ “commute” with coinitial traces \mathbf{t} , giving rise to residuals \mathbf{t}/γ and γ/\mathbf{t} (§4.2). These are used to define causal equivalences $\alpha : \mathbf{t} \simeq \mathbf{u}$ and composite braidings γ_α relating $\mathbf{tgt}(\mathbf{t})$ and $\mathbf{tgt}(\mathbf{u})$ (§4.3).

4.1. Traces

Define $\mathbf{a} : \text{Action}^* \Gamma$ (bold \mathbf{a}) to be a finite sequence of composable actions starting at Γ , where \mathbf{a} and \mathbf{a}' are composable iff $\mathbf{a} \in \text{Action} \Gamma$ and $\mathbf{a}' \in \text{Action} (\Gamma + \mathbf{tgt}(\mathbf{a}))$. $|\mathbf{a}|$ denotes the sum of $|\mathbf{a}|$ for every \mathbf{a} in \mathbf{a} . The empty sequence (nil) at Γ is written ε_Γ ; extension to the left (cons) is written $\mathbf{a} \cdot \mathbf{a}$. A *trace* $\mathbf{t} : P \xrightarrow{\mathbf{a}} R$ (bold \mathbf{t}) is a finite sequence of composable transitions, where \mathbf{t} and \mathbf{u} are composable iff $\text{src}(\mathbf{u}) = \mathbf{tgt}(\mathbf{t})$. The nil trace at P is written ε_P ; cons of $\mathbf{t} : P \xrightarrow{\mathbf{a}} R$ onto $\mathbf{t} : R \xrightarrow{\mathbf{a}} S$ is written $\mathbf{t} \cdot \mathbf{t} : P \xrightarrow{\mathbf{a} \cdot \mathbf{a}} S$.

The renamings $\rho^* \mathbf{a}$ and $\rho^* \mathbf{t}$ of an action and a transition extend to action sequences and traces respectively.

Lemma 11 (Lifting of renamings to action sequences and traces).

Suppose $\rho : \Gamma \longrightarrow \Delta$ and $\mathbf{t} : P \xrightarrow{\mathbf{a}} R$, where $\Gamma \vdash P$, $\Gamma + \Gamma' \vdash R$ and $\mathbf{a} : \text{Action}^* \Gamma$.

$$\begin{array}{ccc} \Gamma \vdash P & \xrightarrow{\mathbf{t}^{\mathbf{a}}} & \Gamma + \Gamma' \vdash R \\ \rho^* \downarrow & & \downarrow (\rho + \Gamma')^* \\ \Delta \vdash \rho^* P & \xrightarrow{(\rho^* \mathbf{t})^{\rho^* \mathbf{a}}} & \Delta + \Gamma' \vdash (\rho + \Gamma')^* R \end{array}$$

Then there exist actions $\rho^* \mathbf{a} : \text{Action}^* \Delta$ and trace $\rho^* \mathbf{t} : \rho^* P \xrightarrow{\rho^* \mathbf{a}} (\rho + \Gamma')^* R$.

Proof. By the following defining equations.

$$\begin{aligned}
\rho^* \varepsilon_\Gamma &= \varepsilon_\Delta & \rho^*(b \cdot a) &= (\rho^*b) \cdot (\rho + 1)^*a \\
\rho^* \varepsilon_P &= \varepsilon_P & \rho^*(c \cdot a) &= (\rho^*c) \cdot \rho^*a \\
& & \rho^*(t^b \cdot t) &= (\rho^*t^b) \cdot (\rho + 1)^*t \\
& & \rho^*(t^c \cdot t) &= (\rho^*t^c) \cdot \rho^*t
\end{aligned}$$

□

4.2. Residuals of traces and braidings

We now develop a minimal generalisation of our system of transitions and braidings sufficient to admit the following notions of residuation:

$$\begin{array}{ccc}
\Gamma \vdash P & \xrightarrow{t} & \Gamma + \Delta \vdash R \\
\gamma \Big| & & \Big| \gamma/t \\
\Gamma \vdash P' & \xrightarrow[t/\gamma]{} & \Gamma + \Delta \vdash R'
\end{array}$$

so that we can accommodate the scenario illustrated earlier in Figure 10. Here $\Delta \in \{0, 1\}$ and γ is a braiding witnessing the cofinality of the target states of an earlier concurrent transition. Recall from Definition 4 that γ relates P and P' either by \times (free braid), \bowtie (bound braid) or $=$ (cofinality “on the nose”); we consider each case and explain how cofinality must be extended to support γ/t . The final definitions of the two residuals are given as the proof of Lemma 13 below.

Case $P \times P'$. Then $P = \text{swap}^*P'$ and $R = (\text{swap} + \Delta)^*R'$ by Lemma 9. If $\Delta = 1$ then the free braid has shifted under a binder and thus $R \not\bowtie R'$. Therefore the first generalisation closes free braids under translations by an arbitrary Δ , allowing them to be preserved by subsequent computation involving bound actions which open up the process term. We define the following relation, noting that $\times = \times_0$.

Definition 5 (Free braid, generalised). For any processes $\Gamma + 2 + \Delta \vdash P, R$ define the symmetric relation $P \bowtie_\Delta R$ as follows. The context Γ is left implicit.

$$P \bowtie_\Delta R \quad \Leftrightarrow \quad P = (\text{swap}_\Gamma + \Delta)^*R$$

Case $P \bowtie P'$. Whereas a free braid inserts a **swap** renaming at the root of P , a bound braid inserts a **swap** under exactly one pair of adjacent binders in P , and thus points to a specific location common to P and P' . When a transition $t : P \xrightarrow{a} R$ is taken, subterms of P may be dropped or duplicated: in particular non-taken branches of choices are discarded, and the bodies of replications are copied into both sides of the resulting parallel compositions. It may therefore not be possible to obtain R' from R by inserting exactly one bound **swap**, since the braid might have been duplicated or thrown away. The second generalisation thus closes bound braids under reflexivity (to permit dropping) and parallel composition (to permit duplication). Figure 12 defines the new relation, also written \bowtie .

$P \bowtie R$

$$\begin{array}{c}
\text{vv-swap}_P \frac{P \bowtie R}{vvP \bowtie vvR} \quad \mathbf{0} \frac{}{\mathbf{0} \bowtie \mathbf{0}} \quad \underline{x}.P \frac{}{\underline{x}.P \bowtie \underline{x}.P} \quad \overline{x}(y).P \frac{}{\overline{x}(y).P \bowtie \overline{x}(y).P} \\
\cdot + Q \frac{P \bowtie R}{P + Q \bowtie R + Q} \quad P + \cdot \frac{Q \bowtie S}{P + Q \bowtie P + S} \quad \cdot \mid \cdot \frac{P \bowtie R \quad Q \bowtie S}{P \mid Q \bowtie R \mid S} \\
\mathbf{v} \cdot \frac{P \bowtie R}{\mathbf{v}P \bowtie \mathbf{v}R} \quad \mathbf{!} \cdot \frac{P \bowtie R}{\mathbf{!}P \bowtie \mathbf{!}R}
\end{array}$$

Fig. 12: Bound braid $P \bowtie R$ that can be dropped or duplicated

Case $P = P'$. The situation is trivial, since t/γ is just t and so γ/t is simply the reflexivity proof that $R = R'$.

The three cases above determine a new braiding relation $\bowtie_{\hat{a}, \Delta}$ which is closed under transitions.

Definition 6 (Braiding, generalised). For any contexts Γ, Δ , any $a, a' \in \text{Action } \Gamma$ and any $\hat{a} : a \sim a'$, define the following symmetric relation $\bowtie_{\hat{a}, \Delta}$ over processes in $\Gamma' + \Delta$, where Γ' is the target context of \hat{a} . There are only two cases rather than three, since the $=$ case is now subsumed by the reflexivity of bound braids.

$$\bowtie_{\hat{a}, \Delta} \stackrel{\text{def}}{=} \begin{cases} \bowtie_{\Delta} & \text{if } \bowtie_{\hat{a}} = \bowtie \\ \bowtie & \text{otherwise} \end{cases}$$

Since $\bowtie = \bowtie_0$, and there is an obvious embedding, via reflexivity, of the old definition of \bowtie (Figure 9) into the new one, there is also an embedding of $\bowtie_{\hat{a}}$ into $\bowtie_{\hat{a}, 0}$.

Lemma 12. $\bowtie_{\hat{a}} \subseteq \bowtie_{\hat{a}, 0}$

The new braidings are sufficiently general to be closed under transitions, so we can go ahead and define the required residuals γ/t and t/γ . We start with the case when γ is a bound braid ϕ . Note that subsuming the $=$ case into the reflexivity of \bowtie does not lose any precision, since for any $t : P \xrightarrow{a} R$ we have $t/P_{\bowtie} = t$ and thus $P_{\bowtie}/t = R_{\bowtie}$.

Theorem 2. Suppose $t : P \xrightarrow{a} R$ and $\phi : P \bowtie P'$. Then there exists a process R' , transition $t/\phi : P' \xrightarrow{a} R'$ and bound braid $\phi/t : R \bowtie R'$.

$$\begin{array}{ccc}
P & \xrightarrow{t} & R \\
\phi \downarrow & & \downarrow \phi/t \\
P' & \xrightarrow{t/\phi} & R'
\end{array}$$

Proof. By the defining equations in Figure 13. Unlike residuals of the form t/t' , the cofinality of t/ϕ and ϕ/t is by construction. P_{\bowtie} denotes the reflexivity proof that $P \bowtie P$.

t/ϕ

$$\begin{aligned}
(\overline{v}v^{\overline{x+1}(0)}t)/vv\text{-swap}_{\text{src}(t)} &= v^{\overline{x}}\overline{v}(\text{swap}^*t) \\
(v^{\overline{x}}\overline{v}t)/vv\text{-swap}_{\text{src}(t)} &= \overline{v}v^{\overline{x+1}(0)}(\text{swap}^*t) \\
(v^c v^{c'}t)/vv\text{-swap}_{\text{src}(t)} &= v^c v^{c'}(\text{swap}^*t) \\
(v^b v^{b'}t)/vv\text{-swap}_{\text{src}(t)} &= v^b v^{b'}(\text{swap}^*t) \\
(\underline{x}.P)/(\underline{x}.\phi) &= \underline{x}.\text{tgt}(\phi) \\
(\overline{x}\langle y \rangle.P)/(\overline{x}\langle y \rangle.\phi) &= \overline{x}\langle y \rangle.\text{tgt}(\phi) \\
(t+Q)/(\phi+Q) &= t/\phi+Q \\
(t+Q)/(P+\psi) &= t+\text{tgt}(\psi) \\
(P+u)/(P+\psi) &= P+u/\psi \\
(P+u)/(\phi+Q) &= \text{tgt}(\phi)+u \\
(t^b|Q)/(\phi|\psi) &= t/\phi^b|\text{tgt}(\psi) \\
(t^c|Q)/(\phi|\psi) &= t/\phi^c|\text{tgt}(\psi) \\
(P|^b u)/(\phi|\psi) &= \text{tgt}(\phi)|^b u/\psi \\
(P|^c u)/(\phi|\psi) &= \text{tgt}(\phi)|^c u/\psi \\
(t|_y u)/(\phi|\psi) &= t/\phi|_y u/\psi \\
(t_y|u)/(\phi|\psi) &= t/\phi_y|u/\psi \\
(t|_v u)/(\phi|\psi) &= t/\phi|_v u/\psi \\
(t_v|u)/(\phi|\psi) &= t/\phi_v|u/\psi \\
(\overline{v}t)/(v\phi) &= \overline{v}t/\phi \\
(v^b t)/(v\phi) &= v^b t/\phi \\
(v^c t)/(v\phi) &= v^c t/\phi \\
(!t)/(!\phi) &= !t/(\phi|!\phi)
\end{aligned}$$

 ϕ/t

$$\begin{aligned}
vv\text{-swap}_{\text{src}(t)}/(\overline{v}v^{\overline{x+1}(0)}t) &= v\text{tgt}(t)_{\times} \\
vv\text{-swap}_{\text{src}(t)}/(v^{\overline{x}}\overline{v}t) &= v(\text{swap}^*\text{tgt}(t))_{\times} \\
vv\text{-swap}_{\text{src}(t)}/(v^c v^{c'}t) &= vv\text{-swap}_{\text{tgt}(t)} \\
vv\text{-swap}_{\text{src}(t)}/(v^b v^{b'}t) &= vv\text{-swap}_{\text{swap}^*(\text{swap}+1)^*\text{swap}^*\text{tgt}(t)} \\
(\underline{x}.\phi)/(\underline{x}.P) &= \phi \\
(\overline{x}\langle y \rangle.\phi)/(\overline{x}\langle y \rangle.P) &= \phi \\
(\phi+Q)/(t+Q) &= \phi/t \\
(P+\psi)/(t+Q) &= \text{tgt}(t)_{\times} \\
(P+\psi)/(P+u) &= \psi/u \\
(\phi+Q)/(P+u) &= \text{tgt}(u)_{\times} \\
(\phi|\psi)/(t^b|Q) &= \phi/t|\text{push}^*\psi \\
(\phi|\psi)/(t^c|Q) &= \phi/t|\psi \\
(\phi|\psi)/(P|^b u) &= \text{push}^*\phi|\psi/u \\
(\phi|\psi)/(P|^c u) &= \phi|\psi/u \\
(\phi|\psi)/(t|_y u) &= (\text{pop } y)^*\phi/t|\psi/u \\
(\phi|\psi)/(t_y|u) &= \phi/t|(\text{pop } y)^*\psi/u \\
(\phi|\psi)/(t|_v u) &= v(\phi/t|\psi/u) \\
(\phi|\psi)/(t_v|u) &= v(\phi/t|\psi/u) \\
(v\phi)/(\overline{v}t) &= \phi/t \\
(v\phi)/(v^b t) &= v\text{swap}^*\phi/t \\
(v\phi)/(v^c t) &= v\phi/t \\
(!\phi)/(!t) &= (\phi|!\phi)/t
\end{aligned}$$

Fig. 13: Residuals of transition t and coinital bound braid ϕ

Figure 14 illustrates Theorem 2 for the cases where ϕ is of the form $vv\text{-swap}_P$, omitting the various renaming lemmas used as type-level coercions.

$$\begin{array}{cc}
\begin{array}{c}
\Gamma \vdash vv(\text{swap}^*P) \xrightarrow{\overline{v}v^{\overline{x+1}(0)}t} \Gamma + 1 \vdash vR \\
\text{vv-swap}_P \Big| \\
\Gamma \vdash vvP \xrightarrow{v^{\overline{x}}\overline{v}(\text{swap}^*t)} \Gamma + 1 \vdash vR
\end{array}
&
\begin{array}{c}
\Gamma \vdash vv(\text{swap}^*P) \xrightarrow{v^c v^{c'}t} \Gamma \vdash vvR \\
\text{vv-swap}_P \Big| \\
\Gamma \vdash vvP \xrightarrow{v^c v^{c'}(\text{swap}^*t)} \Gamma \vdash vv(\text{swap}^*R)
\end{array} \\
\begin{array}{c}
\Gamma \vdash vv(\text{swap}^*P) \xrightarrow{v^{\overline{x}}\overline{v}t} \Gamma + 1 \vdash v(\text{swap}^*R) \\
\text{vv-swap}_P \Big| \\
\Gamma \vdash vvP \xrightarrow{\overline{v}v^{\overline{x+1}(0)}(\text{swap}^*t)} \Gamma + 1 \vdash v(\text{swap}^*R)
\end{array}
&
\begin{array}{c}
\Gamma \vdash vv(\text{swap}^*P) \xrightarrow{v^b v^{b'}t} \Gamma + 1 \vdash vv((\text{swap}+1)^*\text{swap}^*R) \\
\text{vv-swap}_P \Big| \\
\Gamma \vdash vvP \xrightarrow{v^b v^{b'}(\text{swap}^*t)} \Gamma + 1 \vdash vv(\text{swap}^*(\text{swap}+1)^*\text{swap}^*R)
\end{array}
\end{array}$$

Fig. 14: Cofinality of ϕ/t and t/ϕ in the $vv\text{-swap}$ cases

It is then straightforward to extend the bound braid cases t/ϕ and ϕ/t to arbitrary braidings γ and sequences of transitions \mathbf{t} .

Lemma 13 (Residuals of transition t and γ). Suppose $t : P \xrightarrow{a} R$ and $\gamma : P \bowtie_{\hat{a}, \Delta} P'$. Then there exists process R' , action a/γ , transition $t/\gamma : P' \xrightarrow{a/\gamma} R'$ and braiding $\gamma/t : R \bowtie_{\hat{a}, \Delta'} R'$, where $\Delta' = \Delta + |a|$.

$$\begin{array}{ccc} \Gamma + \Delta \vdash P & \xrightarrow{t} & \Gamma + \Delta' \vdash R \\ \gamma \downarrow & & \downarrow \gamma/t \\ \Gamma + \Delta \vdash P' & \xrightarrow{t/\gamma} & \Gamma + \Delta' \vdash R' \end{array}$$

Proof. By the following defining equations, which are given for t/γ and γ/t simultaneously. As before $P_{=}$ denotes the reflexivity proof that $P = P$.

$$(t/\gamma, \gamma/t) = \begin{cases} ((\text{swap} + \Delta)^* t, ((\text{swap} + \Delta')^* R)_{=}) & \text{if } P \bowtie_{\Delta} P' \\ (t/\phi, \phi/t) & \text{if } P \bowtie P' \text{ and } \gamma = \phi \end{cases}$$

The diagram for Lemma 14 is the same as for Lemma 13 but with \mathbf{t} instead of t .

Lemma 14 (Residuals of trace \mathbf{t} and γ).

Suppose $\mathbf{t} : P \xrightarrow{a} R$ and $\gamma : P \bowtie_{\hat{a}, \Delta} P'$. Then there exists process R' , action sequence a/γ , trace $\mathbf{t}/\gamma : P' \xrightarrow{a/\gamma} R'$ and braiding $\gamma/\mathbf{t} : R \bowtie_{\hat{a}, \Delta'} R'$, where $\Delta' = \Delta + |\mathbf{a}|$.

Proof. By the following defining equations.

$$\begin{array}{ccc} \begin{array}{ccc} P & \xrightarrow{\varepsilon_P} & P \\ \gamma \downarrow & & \downarrow \gamma \\ P' & \xrightarrow{\varepsilon_{P'}} & P' \end{array} & \begin{array}{ccccc} P & \xrightarrow{\mathbf{t}} & R & \xrightarrow{\mathbf{t}} & S \\ \gamma \downarrow & & \downarrow \gamma/\mathbf{t} & & \downarrow (\gamma/\mathbf{t})/\mathbf{t} \\ P' & \xrightarrow{\mathbf{t}/\gamma} & R' & \xrightarrow{\mathbf{t}/(\gamma/\mathbf{t})} & S' \end{array} \\ \varepsilon_P/\gamma = \varepsilon_{P'} & (\mathbf{t} \cdot \mathbf{t})/\gamma = \mathbf{t}/\gamma \cdot \mathbf{t}/(\gamma/\mathbf{t}) \\ \gamma/\varepsilon_P = \gamma & \gamma/(\mathbf{t} \cdot \mathbf{t}) = (\gamma/\mathbf{t})/\mathbf{t} \end{array}$$

4.3. Causal equivalence

A causal equivalence $\alpha : \mathbf{t} \simeq \mathbf{u}$ reorders a trace \mathbf{t} into an equal-length, coinital trace \mathbf{u} by permuting concurrent transitions. Meta-variables α, β range over causal equivalences. If $\alpha : \mathbf{t} \simeq \mathbf{u}$ then $\text{tgt}(\mathbf{t})$ and $\text{tgt}(\mathbf{u})$ are related by a unique braiding γ_α .

In what follows, rules which mention a trace of the form $\mathbf{t} \cdot \mathbf{t}$ have an implicit side-condition asserting $\text{tgt}(\mathbf{t}) = \text{src}(\mathbf{t})$, and rules which mention a braiding $\gamma_{t, t'}$ have an implicit side-condition asserting $\mathbf{t} \sim \mathbf{t}'$.

Definition 7. Inductively define the relation \simeq using the rules in Figure 15, where syntactically \simeq has lower priority than \cdot .

$$t \simeq u$$

$$\begin{array}{c}
\mathcal{E}_P \frac{}{\mathcal{E}_P \simeq \mathcal{E}_P} \qquad t \cdot \cdot \frac{t \simeq u}{t \cdot t \simeq t \cdot u} \\
(t \smile t') \cdot t \frac{}{t \cdot t'/t \cdot t \simeq t' \cdot t/t' \cdot t/\gamma_{t,t'}} \qquad \cdot \circ \cdot \frac{t' \simeq u \quad t \simeq t'}{t \simeq u}
\end{array}$$

Fig. 15: Causal equivalence

The \mathcal{E}_P and $t \cdot \cdot$ rules are the congruence cases. The $\cdot \circ \cdot$ rule closes under transitivity, which is a form of vertical composition and which also causes braidings to compose vertically. The transposition rule $(t \smile t') \cdot t$ composes a concurrent pair $t \smile t'$ with a continuation t for t'/t , transporting t through the braiding $\gamma_{t,t'}$ witnessing the cofinality of t and t' to obtain the continuation $t/\gamma_{t,t'}$ for t/t' , as shown in Figure 16.




Fig. 16: Causal equivalence, transposition rule

Theorem 3. \simeq is an equivalence relation.

Proof. Reflexivity is a trivial induction, using the \mathcal{E}_P and $t \cdot \alpha$ rules. Transitivity is immediate from the $\cdot \circ \cdot$ rule. Symmetry is trivial in the \mathcal{E}_P , $t \cdot \alpha$ and $\alpha \circ \beta$ cases. The $(t \smile t') \cdot t$ case requires the symmetry of \smile and that $(t/\gamma)/\gamma = t$. \square

A causal equivalence $\alpha : t \simeq u$ determines a composite braiding relation \mathbf{x}_α which precisely sequences the atomic braidings required to relate $\mathbf{tgt}(t)$ to $\mathbf{tgt}(u)$.

Definition 8 (Braiding for equivalent traces). Inductively define the family \mathbf{x}_α of relations between processes, for any $\alpha : t \simeq u$, using the rules in Figure 17.

$$P \mathbf{x}_\alpha R$$

$$\begin{array}{c}
\frac{}{P \mathbf{x}_{\mathcal{E}_P} P} \qquad \frac{}{P \mathbf{x}_{(t \smile t') \cdot t} R} \quad \gamma_{t,t'}/t : P \mathbf{x}_{\delta, \Delta} R \qquad \frac{P \mathbf{x}_\alpha R}{P \mathbf{x}_{t \cdot \alpha} R} \qquad \frac{P \mathbf{x}_\beta R \quad R \mathbf{x}_\alpha S}{P \mathbf{x}_{\alpha \circ \beta} S}
\end{array}$$

Fig. 17: Braiding relation \mathbf{x}_α relating $\mathbf{tgt}(t)$ and $\mathbf{tgt}(u)$ for any $\alpha : t \simeq u$

As with $\mathbf{x}_{\hat{\alpha}, \Delta}$, the relation \mathbf{x}_α is a singleton, inhabited by a unique path γ_α between $\mathbf{tgt}(\mathbf{t})$ and $\mathbf{tgt}(\mathbf{u})$. (However α itself is not unique, since there are many ways of proving $\mathbf{t} \simeq \mathbf{u}$.) The $P \mathbf{x}_{\varepsilon_P} P$ case is an empty composite braiding. The $P \mathbf{x}_{(t \sim t') \cdot t} R$ case turns an atomic braiding $\gamma_{t, t'}$ into one step of a composite braiding, after transporting it through the continuation \mathbf{t} . The $P \mathbf{x}_{t \cdot \alpha} R$ case simply recognises that $\mathbf{tgt}(\mathbf{t} \cdot \mathbf{t}) = \mathbf{tgt}(\mathbf{t})$. Finally $P \mathbf{x}_{\alpha \circ \beta} R$ is the composition rule, closing under transitivity.

Theorem 4. Suppose $\alpha : \mathbf{t} \simeq \mathbf{u}$. Then there exists a unique $\gamma_\alpha : \mathbf{tgt}(\mathbf{t}) \mathbf{x}_\alpha \mathbf{tgt}(\mathbf{u})$.

Theorem 5. \mathbf{x}_α is a \simeq -indexed family of equivalence relations.

5. Related work

The μs calculus [Hirschhoff 1999] has a similar treatment of de Bruijn indices. Its renaming operators $\langle x \rangle$, ϕ and ψ are effectively our **pop** x , **push** and **swap** renamings, but fused with the \cdot^* operator which applies a renaming to a process. Hirschhoff’s operators are also syntactic forms in the μs calculus, rather than meta-operations, and therefore the operational semantics also includes rules for reducing occurrences of the renaming operators that arise during a process reduction step.

As noted earlier in the paper, our approach to defining causal equivalence of traces is influenced by a line of work stemming from the study of optimal reduction in the λ -calculus [Lévy 1980], via the “proved transition” semantics of CCS [Boudol and Castellani 1989].

Boreale and Sangiorgi [1998] and Degano and Priami [1999] investigate causality in the context of the π -calculus. Similar ideas (from which we also drew inspiration) appear in work on reversible CCS, such as RCCS [Danos and Krivine 2004], and reversible π -calculi, such as $\rho\pi$ [Lanese et al. 2010] and $R\pi$ [Cristescu et al. 2013]). Reversible calculi equip process terms with additional structure to support undoing actions; causal equivalence and permutation of transitions is necessary here to allow undoing actions in a different (sequential) order than they were performed. However, this additional structure changes the metatheory: for example, in $R\pi$ two traces are coinital and cofinal if and only if they are equivalent, which does not hold in our setting. To the best of our knowledge, there is no prior work that presents a proved transition semantics for a “vanilla” π -calculus, rather than an augmented variant.

Another related concept for concurrency calculi, *confluence*, has been studied for CCS [Milner 1980] and for the π -calculus [Philippou and Walker 1997]. A process is confluent if none of its possible actions interfere with each other. Intuitively, this should be the case if the process has only one possible trace modulo causal equivalence. However, to the best of our knowledge, confluence has not been studied using the proved transitions approach and the formal relationship between confluence and causal equivalence is unclear. Our formalisation provides a platform for future study of this matter.

5.1. Mechanised treatments

Formalisations of the π -calculus have been undertaken in several theorem provers used for mechanised metatheory, including Coq, HOL, Isabelle/HOL, Nominal Isabelle, CLF, Abella, and Agda.

HOL Melham [1994] reports on a formalisation of the π -calculus in HOL, using names axiomatised as an unspecified, infinite set, and following Milner et al. [1992] closely. Substitution is parameterised over a choice function specifying how to choose a name fresh for a given set of names, which is used to rename bound names to avoid capture. Aït Mohamed [1995] formalised the π -calculus in HOL using concrete syntax and verified proof rules for early bisimulation checking.

Coq An early mechanisation of residuation theory was Huet’s formalisation in Coq of residuals for λ -calculus [Huet 1994], which also uses de Bruijn indices. Huet’s chief contribution is an inductive definition of residual, a proof that residuals commute with substitution, and a “prism” theorem that generalises Lévy’s cube lemma.

Hirschhoff [1997] formalised the π -calculus in Coq using de Bruijn indices, and verified properties such as congruence and structural equivalence laws of bisimulation. Despeyroux [2000] formalised the π -calculus in Coq using weak higher-order abstract syntax, assuming a decidable type of names, and using two separate transitions, for ordinary, input and output transitions respectively; for input and output transitions the right-hand side is a function of type `name \rightarrow proc`. This formalisation included a simple type system and proof of type soundness. Honsell et al. [2001] formalised the π -calculus in Coq, also using weak higher-order abstract syntax. The type of names `name` is a type parameter assumed to admit decidable equality and freshness (`notin`) relations. Transitions are encoded using two inductive definitions, for free and bound actions, which differ in the type of the third argument (`proc` vs. `name \rightarrow proc`). Numerous results from Milner et al. [1992] are verified, using the *theory of contexts* (whose axioms are assumed in their formalisation, but have been validated semantically by Bucalo et al. [2006]).

Affeldt and Kobayashi [2008] developed a library based on a variant of the π -calculus (with channels typed using Coq types) for representing and reasoning about concurrent processes. Processes are represented using higher-order abstract syntax, and exotic terms are allowed; some lemmas are not formally proved but introduced as axioms with semantic justifications.

Isabelle/HOL Röckl et al. [2001] and Röckl and Hirschhoff [2003] formalised the π -calculus in Isabelle/HOL and verified properties such as adequacy, following the theory of contexts approach to higher-order abstract syntax introduced by Honsell et al. [2001], and using well-formedness predicates to rule out exotic terms. Gay [2001] developed a framework for formalising (linear) type systems for the π -calculus in Isabelle/HOL, using de Bruijn indices for binding syntax and a reduction-style semantics rather than labelled transitions.

Abella Tiu and Miller [2010] encode the syntax and semantics of the π -calculus using the λ -term abstract syntax variant of higher-order abstract syntax; like a number of other approaches they split the transition relation into two relations to handle scope extrusion. Their formalisations employ the meta-logic $\text{FOL}^{\Delta\nabla}$ which forms the basis of the Abella theorem prover, and similar specifications have been used as the basis for verification of properties of the π -calculus in Abella [Baelde et al. 2014].

Accattoli [2012] adapts Huet’s Coq formalisation of residuals from de Bruijn indices to Abella’s higher-order abstract syntax and nominal quantifier ∇ , yielding a significant simplification of Huet’s proof. Accattoli also proves the cube lemma directly, rather than introducing an intermediate prism theorem. It may be that reformalizing our approach using Abella would make it possible to simplify our proof in a similar way.

Nominal Isabelle The Nominal Datatype Package extension to Isabelle/HOL [Urban 2008] supports the Gabbay-Pitts style “nominal” approach to abstract syntax modulo name-binding [Gabbay and Pitts 2002], and has been used in several formalisations. Two early contributions using similar ideas predate its development: Röckl [2001] formalised the syntax of π -calculus and α -equivalence in Isabelle/HOL. Gabbay [2003] described how to use Gabbay-Pitts nominal abstract syntax to represent the π -calculus, without giving a mechanised formalisation or proofs of properties.

Bengtson and Parrow [2009] report on an extensive formalisation in Nominal Isabelle, including inversion principles up to structural congruence, properties of strong and weak bisimulation, and a proof that an axiomatisation of strong late bisimilarity is sound and complete. They use a single inductively-defined transition relation, whose third argument is a sum type allowing either an ordinary process or a residual process with a distinguished bound name.

CLF Cervesato et al. [2002] formalise synchronous and asynchronous versions of π -calculus in the Concurrent Logical Framework (CLF), and Watkins et al. [2008] develop a static type system and operational semantics modeled on that of Gordon and Jeffrey [2003] for checking correspondence properties of protocols specified in the π -calculus. CLF employs higher-order abstract syntax, linearity and a monadic encapsulation of certain linear constructs that can identify objects such as traces up to causal equivalence. Thus, CLF’s π -calculus encodings naturally induce equivalences on traces satisfying commuting conversions among synchronous operations. However, a non-trivial effort appears necessary to compare CLF’s notion of trace equivalence with others, because traces are quotiented by a definitional equality by default and there is no explicit notion of concurrency or residuation.

Agda Orchard and Yoshida [2015] present a translation from a functional language with effects to a π -calculus with session types and verify some type-preservation properties of the translation in Agda.

6. Conclusions and future work

To the best of our knowledge, we are the first to report on a mechanised formalisation of concurrency, residuation and causal equivalence for the π -calculus. We employed de Bruijn indices to represent binders and names. Formalisations of λ -calculi often employ this technique, but to our knowledge only Orchard and Yoshida also employ de Bruijn indices in a mechanised formalisation of π -calculus. Whilst de Bruijn indices incur a certain level of administrative overhead, the use of dependent types helps tame their complexity: many invariants are automatically checked by the type system rather than requiring additional explicit reasoning.

Our work appears to be the first to align the notion of “proved transitions” from Boudol and Castellani’s work on CCS with “transition proofs” in the π -calculus. This hinges on the capability to manipulate and perform induction or recursion over derivations, and means we can leverage dependent typing so that residuation is defined only for concurrent transitions, rather than on all pairs of transitions. It is worth noting that while CLF’s approach to encoding π -calculus automatically yields an equivalence on traces, it is unclear (at least to us) whether this equivalence is similar to the one we propose, or whether such traces can be manipulated explicitly as proof objects if desired.

The most notable aspect of our development is the generalised diamond lemma, which allows causally equivalent traces to have target states which are not equal “on the nose” but only up to a precise braiding which captures how binders were reordered. These braidings are more explicit in a de Bruijn indices setting, since free as well as bound names must be rewired when binders are transposed. Generalised cofinality may be relevant to modelling concurrency in other languages where concurrent transitions have effects which commute only up to some equivalence relation, such as dynamic memory allocation.

6.1. Future work

One possible future direction would be to explore trace structures explicitly quotiented by causal equivalence, such as dependence graphs [Mazurkiewicz 1987], event structures [Boudol and Castellani 1989], or rigid families [Cristescu et al. 2015]. We are also interested in extending our approach to accommodate structural congruences, and in understanding whether ideas from homotopy type theory [Univalent Foundations Program 2013], such as quotients or higher inductive types, could be applied to ease reasoning about π -calculus traces modulo causal equivalence and structural congruence.

An interesting possibility would be to separately formalise the abstract notion of a “residuation system” parameterised on a notion of cofinality. One could then show that the π -calculus (equipped with a particular notion of name binding) admits such a residuation system, with cofinality suitably instantiated. This would shed light on which aspects of concurrency and causality are specific to the choice of name-binding formalism. Potentially this modular approach would also make it easier to study variants of π -calculus where interaction arises from different communication patterns, such as the join-calculus [Fournet and Gonthier 2002] or polyadic π -calculus [Carbone and Maffei 2003]. Again, it might be possible to model concurrency and causality in these settings

independently of the rewiring issues associated with permuting transitions that manipulate scope.

Acknowledgements

We are grateful to our colleagues in the Programming Languages Interest Group at Edinburgh for useful discussions, to Vít Šefl for assistance with the Agda formalisation, and to the anonymous referees for comments on the paper. Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-13-1-3006, and EPSRC, grant number EP/K034413/1. The U.S. Government and University of Edinburgh are authorized to reproduce and distribute reprints for their purposes notwithstanding any copyright notation thereon.

References

- Accattoli, B. (2012). Proof pearl: Abella formalization of λ -calculus cube property. In Hawblitzel, C. and Miller, D., editors, *Certified Programs and Proofs*, volume 7679 of *Lecture Notes in Computer Science*, pages 173–187. Springer Berlin Heidelberg.
- Affeldt, R. and Kobayashi, N. (2008). A coq library for verification of concurrent programs. *Electron. Notes Theor. Comput. Sci.*, 199:17–32.
- Aït Mohamed, O. (1995). Mechanizing a pi-calculus equivalence in hol. In *Proceedings of the 8th International Workshop on Higher Order Logic Theorem Proving and Its Applications*, pages 1–16, London, UK. Springer-Verlag.
- Angiuli, C., Morehouse, E., Licata, D. R., and Harper, R. (2014). Homotopical patch theory. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*, ICFP ’14, pages 243–256, New York, NY, USA. ACM.
- Baelde, D., Chaudhuri, K., Gacek, A., Miller, D., Nadathur, G., Tiu, A., and Wang, Y. (2014). Abella: A system for reasoning about relational specifications. *Journal of Formalized Reasoning*, 7(2).
- Bengtson, J. and Parrow, J. (2009). Formalising the pi-calculus using nominal logic. *Logical Methods in Computer Science*, 5(2:16).
- Boreale, M. and Sangiorgi, D. (1998). A fully abstract semantics for causality in the π -calculus. *Acta Inf.*, 35(5):353–400.
- Boudol, G. and Castellani, I. (1989). Permutation of transitions: An event structure semantics for CCS and SCCS. In Bakker, J., Roever, W.-P., and Rozenberg, G., editors, *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, volume 354 of *LNCS*, pages 411–427. Springer.
- Boudol, G. and Castellani, I. (1991). Flow models of distributed computations: three equivalent semantics for CCS. *Information and Computation*, 114:247–312.
- Bucalo, A., Honsell, F., Miculan, M., Scagnetto, I., and Hofmann, M. (2006). Consistency of the theory of contexts. *J. Funct. Program.*, 16(3):327–372.
- Carbone, M. and Maffei, S. (2003). On the expressive power of polyadic synchronisation in π -calculus. *Nordic J. of Computing*, 10(2):70–98.

- Cervesato, I., Pfenning, F., Walker, D., and Watkins, K. (2002). A concurrent logical framework ii: Examples and applications. Technical Report CMU-CS-02-102, Carnegie Mellon University.
- Cristescu, I., Krivine, J., and Varacca, D. (2013). A compositional semantics for the reversible pi-calculus. In *LICS*, pages 388–397.
- Cristescu, I. D., Krivine, J., and Varacca, D. (2015). *Theoretical Aspects of Computing - ICTAC 2015: 12th International Colloquium, Cali, Colombia, October 29-31, 2015, Proceedings*, chapter Rigid Families for CCS and the π -calculus, pages 223–240. Springer International Publishing.
- Curry, H. B. and Feys, R. (1958). *Combinatory Logic*, volume 1 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, Holland.
- Danos, V. and Krivine, J. (2004). Reversible communicating systems. In Gardner, P. and Yoshida, N., editors, *Concurrency Theory, 15th International Conference, CONCUR '04*, volume 3170 of *LNCS*, pages 292–307. Springer.
- de Bruijn, N. (1972). Lambda-calculus notation with nameless dummies: a tool for automatic formula manipulation with application to the Church-Rosser theorem. *Indagationes Mathematicae*, 34(5):381–392.
- Degano, P. and Priami, C. (1999). Non-interleaving semantics for mobile processes. *Theor. Comput. Sci.*, 216(1-2):237–270.
- Despeyroux, J. (2000). A higher-order specification of the pi-calculus. In *IFIP TCS*, number 1872 in *LNCS*, pages 425–439, London, UK. Springer-Verlag.
- Fournet, C. and Gonthier, G. (2002). The join calculus: A language for distributed mobile programming. In *Applied Semantics: Advanced Lectures*, volume 2395/2002 of *LNCS*, pages 268–332, Berlin / Heidelberg. Springer.
- Gabbay, M. J. (2003). The pi-calculus in FM. In Kamareddine, F., editor, *Thirty-five years of Automating Mathematics*, volume 28 of *Kluwer Applied Logic Series*, pages 247–269. Kluwer.
- Gabbay, M. J. and Pitts, A. M. (2002). A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13:341–363.
- Gay, S. J. (2001). A framework for the formalisation of pi calculus type systems in Isabelle/HOL. In *TPHOLs*, pages 217–232, London, UK. Springer-Verlag.
- Gordon, A. D. and Jeffrey, A. (2003). Typing correspondence assertions for communication protocols. *Theor. Comput. Sci.*, 300(1-3):379–409.
- Hirschhoff, D. (1997). A full formalisation of pi-calculus theory in the calculus of constructions. In *TPHOLs*, pages 153–169.
- Hirschhoff, D. (1999). Handling substitutions explicitly in the pi-calculus. In *Proceedings of the Second International Workshop on Explicit Substitutions: Theory and Applications to Programs and Proofs*.
- Honsell, F., Miculan, M., and Scagnetto, I. (2001). π -calculus in (co)inductive-type theory. *Theor. Comput. Sci.*, 253(2):239–285.
- Huet, G. P. (1994). Residual theory in λ -calculus: A formal development. *Journal of Functional Programming*, 4(3):371–394.
- Lanese, I., Mezzina, C. A., and Stefani, J.-B. (2010). Reversing higher-order pi. In *Concurrency Theory, 21st International Conference, CONCUR '10*, pages 478–493. Springer-Verlag.

- Lévy, J.-J. (1980). Optimal reductions in the lambda-calculus. In Seldin, J. P. and Hindley, J. R., editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*, pages 159–191. Academic Press, New York, NY, USA.
- Mazurkiewicz, A. (1987). Trace theory. In *Advances in Petri Nets 1986, Part II on Petri Nets: Applications and Relationships to Other Models of Concurrency*, number 255 in LNCS, pages 279–324, New York, NY, USA. Springer-Verlag.
- Melham, T. F. (1994). A mechanized theory of the π -calculus in HOL. *Nordic J. of Computing*, 1(1):50–76.
- Milner, R. (1980). *A Calculus of Communicating Systems*. Number 92 in LNCS. Springer-Verlag.
- Milner, R. (1999). *Communicating and mobile systems: the π calculus*. Cambridge University Press, Cambridge, UK.
- Milner, R., Parrow, J., and Walker, D. (1992). A calculus of mobile processes, I and II. *Inf. Comput.*, 100(1):1–77.
- Norell, U. (2009). Dependently typed programming in Agda. In *Advanced Functional Programming*, volume 5832 of LNCS, pages 230–266. Springer.
- Orchard, D. A. and Yoshida, N. (2015). Using session types as an effect system. In *Proceedings Eighth International Workshop on Programming Language Approaches to Concurrency- and Communication-cEntric Software, PLACES 2015, London, UK, 18th April 2015.*, pages 1–13.
- Perera, R. and Cheney, J. (2015). Proof-relevant pi-calculus. In Cervesato, I. and Chaudhuri, K., editors, *Proceedings Tenth International Workshop on Logical Frameworks and Meta Languages: Theory and Practice (LFMTP '15)*, volume 185 of EPTCS, pages 46–70. Open Publishing Association.
- Perera, R., Garg, D., and Cheney, J. (2016). Causally consistent dynamic slicing. In Desharnais, J. and Jagadeesan, R., editors, *Concurrency Theory, 27th International Conference, CONCUR '16, Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Philippou, A. and Walker, D. (1997). On confluence in the pi-calculus. In *Proceedings of the 24th International Colloquium on Automata, Languages and Programming, ICALP '97*, pages 314–324, London, UK, UK. Springer-Verlag.
- Röckl, C. (2001). A first-order syntax for the pi-calculus in isabelle/hol using permutations. *Electr. Notes Theor. Comput. Sci.*, 58(1):1–17.
- Röckl, C. and Hirschhoff, D. (2003). A fully adequate shallow embedding of the π -calculus in isabelle/hol with mechanized syntax analysis. *J. Funct. Program.*, 13(2):415–451.
- Röckl, C., Hirschhoff, D., and Berghofer, S. (2001). Higher-order abstract syntax with induction in isabelle/hol: Formalizing the pi-calculus and mechanizing the theory of contexts. In *FOSSACS, FoSSaCS '01*, pages 364–378, London, UK. Springer-Verlag.
- Sangiorgi, D. and Walker, D. (2001). *The Pi-Calculus - a theory of mobile processes*. Cambridge University Press.
- Stark, E. W. (1989). Concurrent transition systems. *Theoretical Computer Science*, 64(3):221–269.
- Tiu, A. and Miller, D. (2010). Proof search specifications of bisimulation and modal logics for the π -calculus. *ACM Trans. Comput. Logic*, 11(2):13:1–13:35.

- Univalent Foundations Program, T. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. <http://homotopytypetheory.org/book>, Institute for Advanced Study.
- Urban, C. (2008). Nominal techniques in Isabelle/HOL. *J. Autom. Reasoning*, 40(4):327–356.
- Watkins, K., Cervesato, I., Pfenning, F., and Walker, D. (2008). Specifying properties of concurrent computations in CLF. *Electron. Notes Theor. Comput. Sci.*, 199:67–87.

Appendix A. Agda module structure

Figure 18 summarises the module structure of the Agda formalisation.

<i>Utilities</i>	
Ext	Extensions to Agda library, https://github.com/rolyp/agda-stdlib-ext
<i>Core modules</i>	
Action	Actions a
Action.Concur	Concurrent actions $a \smile a'$; residuals a/a'
Action.Seq	Action sequences \mathbf{a}
Braiding.Proc	Bound braids $\phi : P \times P'$
Braiding.Transition	Residuals t/ϕ and ϕ/t
Name	Contexts Γ ; names x
Proc	Processes P
ProofRelevantPi	Include everything; compile to build project
ProofRelevantPiCommon	Common imports from standard library
Ren	Renamings $\rho : \Gamma \longrightarrow \Gamma'$
Ren.Properties	Additional properties relating to renamings
Transition	Transitions $t : P \xrightarrow{a} R$
Transition.Concur	Concurrent transitions $t \smile t'$; residuals t/t'
Transition.Concur.Cofinal	Cofinality witnesses γ
Transition.Concur.Cofinal.Transition	Residuals t/γ and γ/t
Transition.Seq	Transition sequences
Transition.Seq.Cofinal	Residuals t/γ and γ/t ; permutation equivalence $\alpha : t \simeq u$
Transition.Seq.Cofinal.Cofinal	Proof that t/γ and γ/t are (heterogeneously) cofinal
<i>Common sub-modules</i>	
.Ren	Renaming lifted to entity defined in parent module

Fig. 18: Module overview, release 0.3